arXiv:1510.05616v1 [math.NA] 14 Oct 2015

A FAST ALGORITHM FOR SIMULATING MULTIPHASE FLOWS THROUGH PERIODIC GEOMETRIES OF ARBITRARY SHAPE

GARY R. MARPLE*, ALEX BARNETT†, ADRIANNA GILLMAN‡, AND SHRAVAN VEERAPANENI§

Abstract. This paper presents a new boundary integral equation (BIE) method for simulating particulate and multiphase flows through periodic channels of arbitrary smooth shape in two dimensions. The authors consider a particular system—multiple vesicles suspended in a periodic channel of arbitrary shape—to describe the numerical method and test its performance. Rather than relying on the periodic Green's function as classical BIE methods do, the method combines the free-space Green's function with a small auxiliary basis, and imposes periodicity as an extra linear condition. As a result, we can exploit existing free-space solver libraries, quadratures, and fast algorithms, and handle a large number of vesicles in a geometrically complex channel. Spectral accuracy in space is achieved using the periodic trapezoid rule and product quadratures, while a first-order semi-implicit scheme evolves particles by treating the vesicle-channel interactions explicitly. New constraint-correction formulas are introduced that preserve reduced areas of vesicles, independent of the number of time steps taken. By using two types of fast algorithms, (i) the fast multipole method (FMM) for the computation of the vesicle-vesicle and the vesicle-channel hydrodynamic interaction, and (ii) a fast direct solver for the BIE on the fixed channel geometry, the computational cost is reduced to $\mathcal{O}(N)$ per time step where N is the spatial discretization size. Moreover, the direct solver inverts the wall BIE operator at $t = 0$, stores its compressed representation and applies it at every time step to evolve the vesicle positions, leading to dramatic cost savings compared to classical approaches. Numerical experiments illustrate that a simulation with $N = 128{,}000$ can be evolved in less than a minute per time step on a laptop.

Key words. Stokes flow, periodic geometry, spectral methods, boundary integral equations, fast direct solvers

1. Introduction. Suspensions of rigid and/or deformable particles in viscous fluids flowing through confined geometries are ubiquitous in natural and engineering systems. Examples include drop, bubble, vesicle, swimmer, and red blood cell (RBC) suspensions. Understanding the spatial distribution of such particles in confined flows is crucial in a wide range of applications including targeted drug delivery [30], enhanced oil recovery [44], and microfluidics for cell sorting and separation [31]. In several of these applications, the long-time behavior of the suspension is sought. For example: *What is the optimal size and shape of targeted drug carriers that maximizes their ability to reach the vascular walls escaping from flowing RBCs* [30, 15]? *What is the optimal design of a microfluidic device that differentially separates circulating tumor cells from blood cells* [49]? More generally, one is interested in estimating the rheological properties of a given particulate suspension in an applied flow, electric, or magnetic fields. A common mathematical construct that is employed in such a scenario is the periodicity of flow at the inlet and the outlet. Therefore, the natural computational problem that arises is to solve for the transient dynamics of a particulate flow through a confined periodic channel driven either by pressure difference or other stimuli.

In this work, we consider periodization algorithms for *vesicle suspensions* in confined flows. Vesicles—often considered as mimics for biological cells, especially RBCs [54]—are comprised of bilipid membranes enclosing a viscous fluid and their diameter is typically less than $10\mu m$. The membrane mechanics is modeled by the Helfrich energy [26] combined with a local inextensibility constraint. At the length scale of the vesicles, the Reynolds number is extremely small, therefore, the Stokes equations are employed to model the fluid interior and exterior to the vesicles. The suspension dynamics of this system is governed by the nonlinear membrane forces, the vesicle-vesicle and vesicle-channel non-local hydrodynamic interactions, and the applied flow boundary conditions.

Pioneered by Youngren and Acrivos [32, 33], BIE methods are widely used for particulate and other interfacial flows [47]. Their advantages over grid- and mesh-based methods are well-known: reduction in dimensionality, ease of achieving high-order accuracy, and availability of highly scalable fast algorithms. The classical approach for incorporating periodic boundary conditions within the BIE framework is to replace the free-space Green's function with one that satisfies the periodicity condition. This can be expressed as an infinite sum of source images. The double-layer potential defined on an open curve Γ

*Department of Mathematics, University of Michigan, Ann Arbor, MI, 48109, USA. *email:* gmarple@umich.edu.

†Department of Mathematics, Dartmouth College, Hanover, NH, 03755, USA. *email:* ahb@math.dartmouth.edu.

‡Computational and Applied Mathematics, Rice University, Houston, TX, 77005, USA. *email:* adrianna.gillman@rice.edu.

§Department of Mathematics, University of Michigan, Ann Arbor, MI, 48109, USA. *email:* shravan@umich.edu.

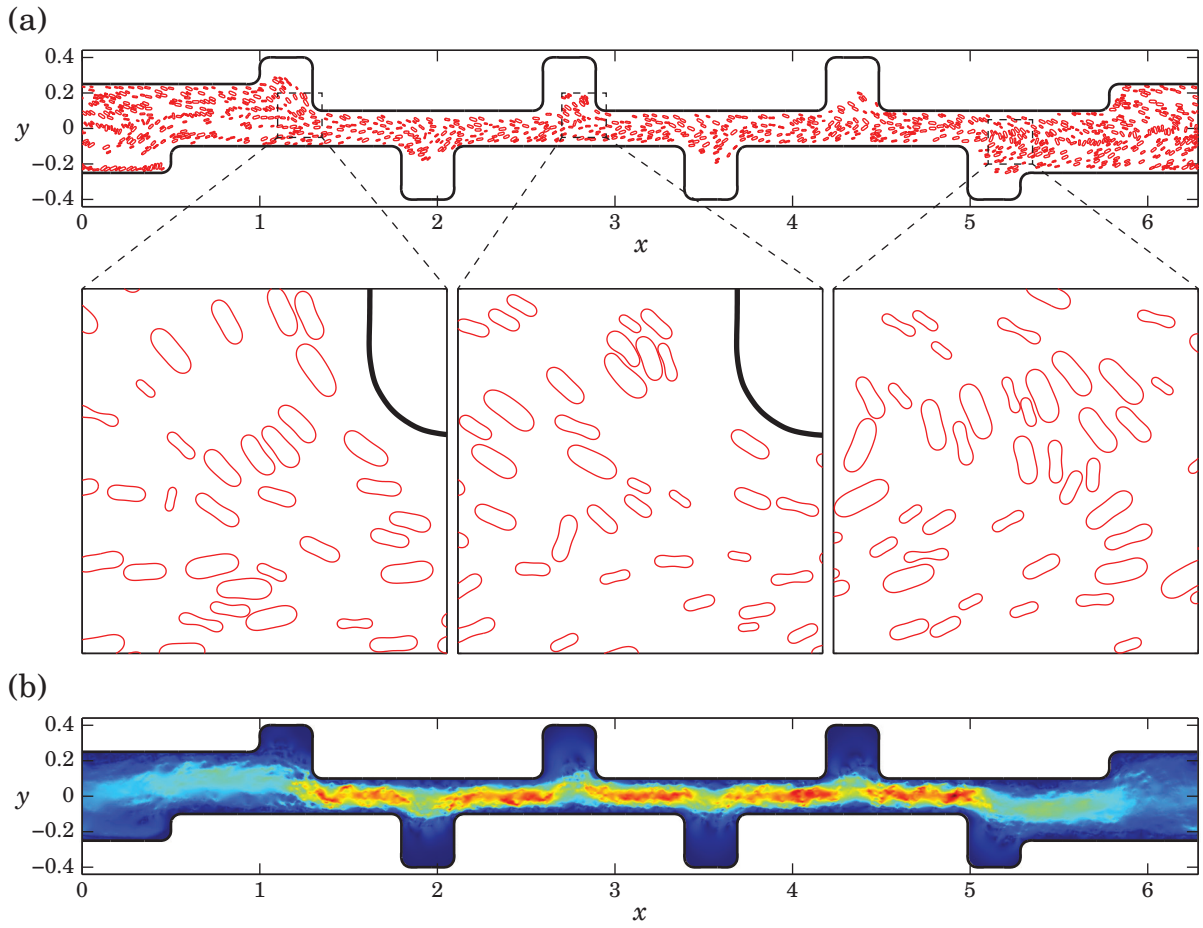

FIG. 1.1. *(a) Snapshot from a simulation of 1,005 vesicles flowing through an arbitrary-shaped periodic channel. We used 64 discretization points per vesicle and 32,000 points each for the top and bottom walls. The vesicle-vesicle and vesicle-channel hydrodynamic interactions are computed via the Stokes FMM, and the new fast direct solver (Sec. 4) is used to solve the channel BIEs. We used the close evaluation scheme of [6] for the vesicle-to-vesicle and vesicle-to-channel interactions (but not for channel-to-vesicle interactions due to its inapplicability). This simulation took 52 seconds per time step on a laptop with a 2.4 GHz dual-core Intel Core i5 processor and 8 GB of RAM. (b) Plot of the velocity magnitude (red indicates high and blue indicates low) corresponding to the disturbance field generated by the vesicles (obtained by subtracting the pressure-driven "empty pipe" flow from the total velocity field).*

which defines one period of a channel wall with lattice vector $\mathbf{d}$, for instance, can be written as

$$\mathbf{u}(\mathbf{x}) = \sum_{n\in\mathbb{Z}} \int_{\Gamma} D(\mathbf{x}, \mathbf{y} + n\mathbf{d})\, \boldsymbol{\tau}(\mathbf{y})\, ds_{\mathbf{y}}, \tag{1.1}$$

where $\mathbf{u}$ is the fluid velocity, D is the Stokes free-space double-layer kernel, and $\boldsymbol{\tau}$ is the density function defined on Γ. Classical algorithms, such as the *Ewald summation* [16, 25, 46] for accelerating the $N-$body calculation that arises from discretizing (1.1), use a partition of unity to split the discrete sum into rapidly converging sums for the *nearby* and *distant* interactions, handling the latter in the spectral domain. The local interactions are $\mathcal{O}(N)$ in number by construction, leaving distant interactions which can be evaluated accurately at the N targets by combining local interpolations onto a regular grid with the fast Fourier transform (FFT), a technique named particle-mesh Ewald [14]. Recently, an accurate variant called spectral Ewald has been developed for particulate flows [39, 1, 55] in periodic geometries.

Although such FFT-based methods are widely used, owing to their ease of implementation, they suffer from several drawbacks. Firstly, the FFT introduces a $\mathcal{O}(N \log N)$ complexity, and although the constants are rather small for FFT methods, the scalability of communication costs on multicore architectures is suboptimal (see [19] for a detailed discussion). Secondly, the lack of spatial adaptivity

makes them somewhat impractical for problems with multi-scale physics. Finally, the "gridding" required is expensive, becoming even more so in three dimensions. One of the main goals of this paper is to introduce a simple alternative algorithm that is $\mathcal{O}(N)$, and, since it exploits existing fast algorithms, overcomes many of these limitations which can be quite restrictive for constrained geometries that have local features (see Fig. 1.1).

Synopsis of the new method. The proposed periodizing integral equation formulation is based on the ideas introduced in Barnett–Greengard [5] for quasi-periodic scattering problems. It uses direct free-space summation for the nearest-neighbor periodic images, whereas the flow field due to the distant images is captured using an auxiliary basis comprised of a small number of "proxy" sources. Periodic boundary conditions are imposed in an extended linear system (ELS) that determines both the wall layer densities (to enforce no-slip boundary condition on the channel) and the proxy source strengths. Although one block of this ELS is rectangular and ill-conditioned, its pseudo-inverse is rapid to compute, allowing accuracy close to machine precision [12, 41]. The disturbance velocities and the hydrodynamic stresses due to the presence of vesicles enter the right-hand side of the ELS so that the combined flow field is periodic from channel inlet to outlet and vanishes on the walls. The proposed integral formulation is versatile in handling the imposed flow boundary conditions: applied pressure-drop across the channel, or imposed slip on the channel (e.g., to model electroosmotic flows), simply modify the right-hand side of the ELS. The scheme can handle various dimensions of periodicity and easily extends to 3D [41].

The proposed particulate flow solver is based on the work of Veerapaneni et al. [53]. It employs a semi-implicit time-stepping scheme to overcome the numerical stiffness associated with the integro-differential equations governing the vesicle evolution. A spectral (Fourier) basis is used to represent the vesicle and channel boundaries. The required spatial derivatives are computed via spectral differentiation and the singular integrals are also computed with spectral accuracy using the product quadrature rule given in Kress [34]. The FMM is used to accelerate the computation of the vesicle-vesicle hydrodynamic interactions. A simple correction term is introduced in the local inextensibility constraint applied at every time step. It eliminates error accumulation over long time-periods that usually leads to numerical instabilities.

Since the channel geometry remains fixed, its BIE linear system may be inverted once and for all. We use a direct solver to precompute its inverse; the channel wall densities can then be determined at each time step for the small cost of a matrix-vector multiply. Since the matrix associated with wall-wall interactions has off-diagonal low-rank structure, the use of a hierarchical fast direct solver reduces the cost of both the precomputation, and of a solve with each new right-hand side to $\mathcal{O}(N)$, where N is the number of points on the channel walls. Crucially, the cost involved with each new solve is very small when compared against the standard combination of an iterative solver and an FMM: for example, Table 5.3 shows that, in our setting, the former is *three orders of magnitude* faster than the latter. The fast direct solver for the ELS is a Stokes version of the periodic Helmholtz scattering solver of Gillman–Barnett [20], with the extra complication that the channel walls form a continuous interface (in [20] it was possible to isolate the obstacles). The continuous interface demands a new compression scheme, presented in Sec. 4. While fast direct solvers have received much attention recently, the authors believe that this work is the first to apply them to particulate flows.

Advantages. The conspicuous advantage of our method is the simplicity that comes from the use of free-space kernels (as opposed to lattice sums or particle-mesh Ewald) and a pressure drop condition that is applied directly. More importantly, this feature allows us to use state-of-the-art high-order quadratures for the singular [36] and nearly-singular integrals [27, 6] and open-source FMM implementations [22]. The periodization scheme itself is spectrally accurate in terms of the number of proxy points. Numerical experiments illustrate that the same number of proxy points are required to achieve a specified accuracy independent of the complexity of the geometry.

Limitations. In this paper, we restrict our attention to two-dimensional problems. Although the periodization scheme itself is straightforward to extend to three dimensions, other components of the numerical method, such as the quadratures and the direct solver for surfaces in three dimensions, require more work. We assume that there are no holes in the given periodic domain and that the viscosities of the fluids interior and exterior to vesicle membranes are the same. Both of these assumptions are merely for simplicity of exposition (e.g., see [48] to relax the latter assumption and the completed double-layer formulation [45] for the former).

Our recent algorithm to evaluate the nearly-singular integrals with spectral accuracy [6] is applied to evaluate the vesicle-vesicle hydrodynamic interactions when they are close to each other. However, we do not apply any close evaluation schemes for the channel-to-vesicle interactions ([6] cannot be applied directly to this setting). This limitation prohibits us from performing simulations of tightly-packed suspensions. One possible remedy is to switch the discretization of the channel from a global basis (Fourier) to local chunk-based schemes, for which close evaluation schemes for Stokes potentials already exist [43].

1.1. Related Work. There is a large body of literature on BIE methods for periodic Stokesian flows of rigid and deformable bodies. A few examples include [60, 38, 17, 23] for fixed particles as in a porous medium, [42, 61, 18, 58, 1] for particulate suspensions, and particularly [59, 51] for vesicle flows. However, almost all studies focus on flows through either simple geometries, such as flat channels or cylinders, or without any constraining walls (i.e., one-, two-, or three-periodic systems in free-space). The work of Greengard–Kropinski [23] uses an intrinsically two-dimensional complex-variable formulation to periodize a FMM-based solver for a fixed doubly-periodic geometry in $\mathcal{O}(N)$ cost. However, the imposition of pressure-drop conditions in their scheme is a subtle matter involving non-convergent lattice sums. In contrast, in the present work, such conditions are applied simply and directly and the cost remains $\mathcal{O}(N)$. The only work that we are aware of where vesicles inhabit an arbitrary periodic geometry is that by Zhao et al [58] where capsules flowing through deformed cylinders were simulated. The constraining geometry is embedded in a box and a Green's function that satisfies periodic conditions on the box is used. One of the drawbacks of this approach is that a pressure drop cannot be imposed directly but is determined from the mean flow. Furthermore, a large amount of auxiliary data might be introduced because of the embedding in the case of geometries that have multi-scale spatial features. The authors would like to point out that to date, there are no reported results on particulate flows through complex periodic geometries such as those shown in Fig. 1.1.

Fast direct solvers, such as $\mathcal{H}$-matrix, HBS, HSS, and HODLR ([57, 50, 56, 8, 9, 21, 2]) which utilized hierarchical low-rank compression of off-diagonal blocks, are naturally applicable to solving the linear systems arising from the discretization of boundary integral equations, thanks to the smoothly decaying property of Green's functions for far interactions. Since the construction of the fast direct solver dominates the computational cost, it is not beneficial to build a solver for the evolving geometries. Instead, the proposed method constructs a fast direct solver for the fixed constrained walls, then reuses this precomputed solver at each time step to evolve the vesicles. The computational cost for both steps scales linearly with the number of discretization points. The cost for each time step is much reduced compared to an iterative FMM-based solve. While this manuscript describes an HBS solver (see [21] or Section 4) for simplicity of presentation, alternative $\mathcal{O}(N)$ inversion techniques can be seamlessly substituted in.

1.2. Outline. This manuscript begins by describing the periodic solver for the steady Stokes equation in a channel given velocity boundary conditions (Section 2). The numerical scheme for the evolution of vesicles and the treatment of the coupling between the vesicle and channel BIEs is presented in Section 3. The fast ELS solver for finding the channel densities and the proxy strengths is presented in Section 4. The accuracy, stability, and computational complexity of the method will be illustrated via numerical experiments in Section 5. Finally, the manuscript concludes with a summary and statement of future work in Section 6.

2. Periodization scheme.

2.1. Preliminaries: Stokes potentials and the non-periodic BVP. We first define the standard kernels and boundary integral operators used [37] (for our 2D case see [29, Sec. 2.2, 2.3]). Let $\mu > 0$ be the fluid viscosity, a scalar constant. Let $\mathbf{u}(\mathbf{x}) = (u_1(\mathbf{x}), u_2(\mathbf{x}))$ be the velocity field and $p(\mathbf{x})$ the scalar pressure field for $\mathbf{x} = (x_1, x_2) \in \mathbb{R}^2$. The pair $(\mathbf{u}, p)$ is a solution to the Stokes equations if

$$-\mu\Delta\mathbf{u} + \nabla p = 0 \tag{2.1}$$

$$\nabla \cdot \mathbf{u} = 0. \tag{2.2}$$

These express force balance and incompressibility, respectively.

The Stokes single-layer kernel (stokeslet) from source point $\mathbf{y}$ to target point $\mathbf{x}$ has tensor components

$$S_{ij}(\mathbf{x},\mathbf{y}) = \frac{1}{4\pi\mu}\left(\delta_{ij}\log\frac{1}{r} + \frac{r_i r_j}{r^2}\right), \quad i,j=1,2, \tag{2.3}$$

where $\mathbf{r} := \mathbf{x}-\mathbf{y}$, $r := \|\mathbf{r}\|$, and δ_{ij} is the Kronecker delta. Given a density (vector function) $\boldsymbol{\tau}$ on a source curve Γ, the single-layer representation for velocity is then $\mathbf{u} = \mathcal{S}_\Gamma\boldsymbol{\tau}$, i.e.,

$$\mathbf{u}(\mathbf{x}) = (\mathcal{S}_\Gamma\boldsymbol{\tau})(\mathbf{x}) := \int_\Gamma S(\mathbf{x},\mathbf{y})\boldsymbol{\tau}(\mathbf{y})ds_{\mathbf{y}}. \tag{2.4}$$

The associated pressure function is

$$p(\mathbf{x}) = (\mathcal{Q}_\Gamma\boldsymbol{\tau})(\mathbf{x}) := \int_\Gamma Q(\mathbf{x},\mathbf{y})\boldsymbol{\tau}(\mathbf{y})ds_{\mathbf{y}} \quad \text{where } Q_j(\mathbf{x},\mathbf{y}) = \frac{1}{2\pi}\frac{r_j}{r^2}\ , \qquad j=1,2. \tag{2.5}$$

For the double-layer velocity representation $\mathbf{u} = \mathcal{D}_\Gamma\boldsymbol{\tau}$, we have, using $\mathbf{n}^{\mathbf{y}}$ the surface normal at each point on the source curve Γ,

$$\mathbf{u}(\mathbf{x}) = (\mathcal{D}_\Gamma\boldsymbol{\tau})(\mathbf{x}) := \int_\Gamma D(\mathbf{x},\mathbf{y})\boldsymbol{\tau}(\mathbf{y})ds_{\mathbf{y}} \quad \text{where } D_{ij}(\mathbf{x},\mathbf{y}) = \frac{1}{\pi}\frac{r_i r_j}{r^2}\frac{\mathbf{r}\cdot\mathbf{n}^{\mathbf{y}}}{r^2}\ , \qquad i,j=1,2. \tag{2.6}$$

We write its associated pressure function as

$$p(\mathbf{x}) = (\mathcal{P}_\Gamma\boldsymbol{\tau})(\mathbf{x}) := \int_\Gamma P(\mathbf{x},\mathbf{y})\boldsymbol{\tau}(\mathbf{y})ds_{\mathbf{y}} \quad \text{where } P_j(\mathbf{x},\mathbf{y}) = \frac{\mu}{\pi}\left(-\frac{\mathbf{n}^{\mathbf{y}}_j}{r^2} + 2\mathbf{r}\cdot\mathbf{n}^{\mathbf{y}}\frac{r_j}{r^4}\right)\ , \qquad j=1,2. \tag{2.7}$$

We use the notation $D_{\Gamma',\Gamma}$ to indicate the double-layer boundary integral operator from source curve Γ to target Γ', i.e., $D_{\Gamma',\Gamma}\boldsymbol{\tau} = (\mathcal{D}_\Gamma\boldsymbol{\tau})|_{\Gamma'}$. If the target and source curves are the same ($\Gamma'=\Gamma$) then $D_{\Gamma,\Gamma}$ is to be taken in the principal value sense and has a smooth kernel for smooth Γ. We have, for Γ a C^2-smooth curve and any $\boldsymbol{\tau}\in C(\Gamma)$, the jump relation

$$\lim_{h\to 0^+}(\mathcal{D}_\Gamma\boldsymbol{\tau})(\mathbf{x}-h\mathbf{n}^{\mathbf{x}}) = (-\tfrac{1}{2}I + D_{\Gamma,\Gamma})\boldsymbol{\tau} \tag{2.8}$$

for the interior limit of velocity. Here, I is the 2×2 identity tensor. The non-periodic prototype BVP that we will need to solve is that the pair $(\mathbf{u},p)$ satisfies the Stokes equations in a bounded domain Ω for given velocity (Dirichlet) data $\mathbf{u}=\mathbf{v}$ on its boundary $\partial\Omega$. To solve this problem, we insert the double-layer representation $\mathbf{u} = \mathcal{D}_{\partial\Omega}\boldsymbol{\tau}$ into (2.8) to get the 2nd-kind boundary integral equation (BIE) on $\partial\Omega$:

$$(-\tfrac{1}{2}I + D_{\partial\Omega,\partial\Omega})\boldsymbol{\tau} = \mathbf{v}\ . \tag{2.9}$$

This BVP, and the resulting BIE, has one consistency condition, $\int_{\partial\Omega}\mathbf{v}\cdot\mathbf{n}ds_{\mathbf{y}} = 0$, and null-space of dimension one corresponding to adding a constant to p [29].

Since we will also need to impose *traction* (Neumann) matching conditions, we need the traction on a target curve due to the above representations. Given a function pair $(\mathbf{u},p)$, the Cauchy stress tensor at any point has entries

$$\sigma_{ij}(\mathbf{u},p) := -\delta_{ij}p + \mu(\partial_i u_j + \partial_j u_i)\ , \quad i,j=1,2. \tag{2.10}$$

The hydrodynamic traction of this pair, i.e., the force vector per unit length applied to the fluid at a surface point with outward unit normal $\mathbf{n}$, has components

$$T_i(\mathbf{u},p) := \sigma_{ij}(\mathbf{u},p)n_j = -pn_i + \mu(\partial_i u_j + \partial_j u_i)n_j\ , \quad i=1,2, \tag{2.11}$$

where summation over j is implied. Applying (2.11) to the $(\mathbf{u}(\mathbf{x}),p(\mathbf{x}))$ pair, due to the single-layer velocity (2.3) and pressure (2.5) kernel (with fixed source point $\mathbf{y}$), gives the single-layer traction kernel

$$K_{ik}(\mathbf{x},\mathbf{y}) = \sigma_{ij}(S_{jk}(\cdot,\mathbf{y}),Q_k(\cdot,\mathbf{y}))(\mathbf{x})\mathbf{n}^{\mathbf{x}}_j = -\frac{1}{\pi}\frac{r_i r_k}{r^2}\frac{\mathbf{r}\cdot\mathbf{n}^{\mathbf{x}}}{r^2}\ , \qquad i,k=1,2, \tag{2.12}$$

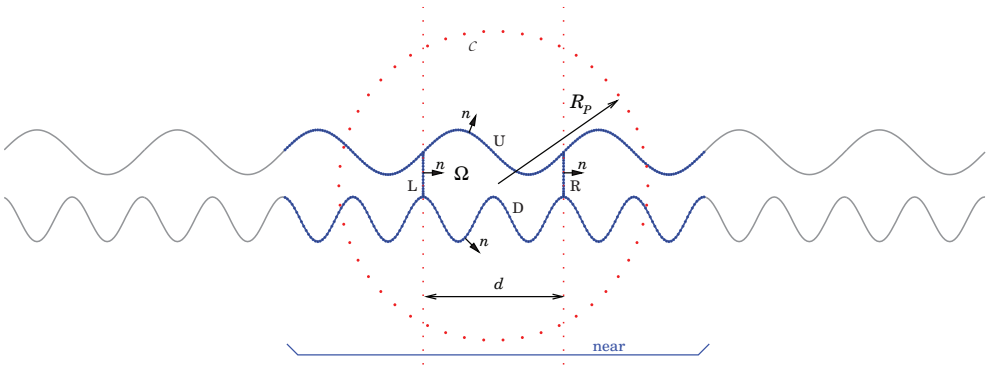

FIG. 2.1. *Geometry for periodization scheme of Section 2. The grey shows the infinite periodic pipe, and the blue dots the quadrature nodes for the central domain and its near neighbors. The central domain Ω is bounded by $\Gamma = U \cup D$, and side walls L and R, and has the normal senses shown. The proxy points (red) lie on circle $\mathcal{C}$ of radius R_p.*

which we abbreviate by K. Likewise, applying (2.11) to the double-layer pair (2.6) and (2.7) gives, after a somewhat involved calculation (eg [40, (5.27)]), the double-layer traction kernel tensor

$$\begin{aligned} T_{ik}(\mathbf{x},\mathbf{y}) &= \sigma_{ij}(D_{jk}(\cdot,\mathbf{y}), P_k(\cdot,\mathbf{y}))(\mathbf{x})\mathbf{n}^{\mathbf{x}}_j \,, \qquad i,k=1,2 \\ &= \frac{\mu}{\pi}\left[\left(\frac{\mathbf{n}^{\mathbf{y}}\cdot\mathbf{n}^{\mathbf{x}}}{r^2} - 8d_{\mathbf{x}}d_{\mathbf{y}}\right)\frac{r_i r_k}{r^2} + d_{\mathbf{x}}d_{\mathbf{y}}\delta_{ik} + \frac{\mathbf{n}^{\mathbf{x}}_i\mathbf{n}^{\mathbf{y}}_k}{r^2} + d_{\mathbf{x}}\frac{r_k\mathbf{n}^{\mathbf{y}}_i}{r^2} + d_{\mathbf{y}}\frac{r_i\mathbf{n}^{\mathbf{x}}_k}{r^2}\right], \end{aligned} \tag{2.13}$$

where we defined the target and source "dipole functions"

$$d_{\mathbf{x}} = d_{\mathbf{x}}(\mathbf{x},\mathbf{y}) := (\mathbf{r}\cdot\mathbf{n}^{\mathbf{y}})/r^2 \,, \qquad\qquad d_{\mathbf{y}} = d_{\mathbf{y}}(\mathbf{x},\mathbf{y}) := (\mathbf{r}\cdot\mathbf{n}^{\mathbf{x}})/r^2,$$

respectively. The use of the symbol T to mean the traction operator and the double-layer traction kernel will be clear by context. The hypersingular boundary integral operator for the traction of the double-layer from source curve Γ to target Γ' we call $T_{\Gamma',\Gamma}$. To clarify,

$$(T_{\Gamma',\Gamma}\boldsymbol{\tau})(\mathbf{x}) = T(\mathcal{D}_\Gamma\boldsymbol{\tau}, \mathcal{P}_\Gamma\boldsymbol{\tau})(\mathbf{x}) = \int_\Gamma T(\mathbf{x},\mathbf{y})\boldsymbol{\tau}(\mathbf{y})ds_{\mathbf{y}} \,, \qquad \mathbf{x}\in\Gamma',$$

where in the final expression, the kernel T has tensor components (2.13).

2.2. Dirichlet problem in a driven periodic pipe. We consider a single unit cell Ω confined by one period of the walls U above and D below. The full periodic pipe domain is then $\Omega_\Lambda := \{\mathbf{x}\in\mathbb{R}^2 : \mathbf{x}+n\mathbf{d}\in\Omega, n\in\mathbb{Z}\}$, where $\mathbf{d}=(d,0)$ is the lattice vector with period d. See Fig. 2.1.

It is conceptually simplest to begin with the following strictly-periodic "empty pipe" problem. Given periodic velocity (Dirichet) data $\mathbf{v}_U$ and $\mathbf{v}_D$ on the up and down walls, find a solution $(\mathbf{u},p)$ in Ω_Λ that is periodic up to a constant pressure driving per period, i.e.,

$$(\mathbf{u},p) \qquad \text{Stokes in } \Omega_\Lambda \tag{2.14}$$

$$\mathbf{u} = \mathbf{v}_U \text{ on } U \tag{2.15}$$

$$\mathbf{u} = \mathbf{v}_D \text{ on } D \tag{2.16}$$

$$\mathbf{u}(\mathbf{x}+\mathbf{d}) - \mathbf{u}(\mathbf{x}) = 0, \quad \mathbf{x}\in\Omega_\Lambda \tag{2.17}$$

$$p(\mathbf{x}+\mathbf{d}) - p(\mathbf{x}) = p_{\text{drive}}, \quad \mathbf{x}\in\Omega_\Lambda. \tag{2.18}$$

The consistency condition on the data is $\int_U \mathbf{v}_U\cdot\mathbf{n}ds_{\mathbf{y}} + \int_D \mathbf{v}_D\cdot\mathbf{n}ds_{\mathbf{y}} = 0$, and the nullity 1, as in the non-periodic case. In our application, the data $\mathbf{v}_D$, $\mathbf{v}_U$ will be (minus) the flow velocity induced by a periodized set of vesicles inside Ω_Λ.

A standard approach for solving this BVP in the strictly-periodic case $p_{\text{drive}}=0$ would be to sum the double-layer kernel over all periodic copies in order to obtain the periodized version of the kernel:

$$D^P(\mathbf{x},\mathbf{y}) = \sum_{n\in\mathbb{Z}} D(\mathbf{x},\mathbf{y}+n\mathbf{d}). \tag{2.19}$$

The representation is then, using $\Gamma = U \cup D$ to indicate the one period of the upper and lower walls,

$$\mathbf{u}(\mathbf{x}) = (\mathcal{D}^P_\Gamma \boldsymbol{\tau})(\mathbf{x}) = \int_U D^P(\mathbf{x},\mathbf{y})\boldsymbol{\tau}_U(\mathbf{y})\,ds_\mathbf{y} + \int_D D^P(\mathbf{x},\mathbf{y})\boldsymbol{\tau}_D(\mathbf{y})\,ds_\mathbf{y}. \tag{2.20}$$

By analogy with (2.9), the density $\boldsymbol{\tau} = [\boldsymbol{\tau}_U; \boldsymbol{\tau}_D]$ could be determined by solving the 2nd-kind integral equation

$$\left(-\tfrac{1}{2}I + \mathcal{D}^P_{\Gamma,\Gamma}\right)\boldsymbol{\tau} = \mathbf{v} \tag{2.21}$$

with $\mathbf{v} = [\mathbf{v}_U; \mathbf{v}_D]$.

REMARK 2.1. *If $p_{drive} \neq 0$, this approach could also be used after substracting from $\mathbf{v}$ velocity data from the Poiseuille flow $\mathbf{u}(\mathbf{x}) = (\frac{1}{2}\alpha x_2^2, 0)$, $p(\mathbf{x}) = \alpha\mu x_1$, with $\alpha = p_{drive}/(\mu d)$. The result is a strictly-periodic BVP with $p_{drive} = 0$ and modified data $\mathbf{v}$. However, we will find the following approach much more convenient.*

Instead, we reformulate the BVP on the single unit cell Ω, introducing a left side wall L and right side wall $R = L + \mathbf{d}$ (Fig. 2.1). Note that, given a periodic pipe Ω_Λ, the choice of where to place the wall to subdivide the unit cell is arbitrary. We choose them to be vertical for convenience. Furthermore we relax the periodicity condition on the wall velocity data $\mathbf{v}_U$, $\mathbf{v}_D$, and impose between L and R periodicity conditions for velocity and traction with given arbitrary mismatch $\mathbf{g}_u$, $\mathbf{g}_T$, that we call the "discrepancies" [5]. Thus,

$$(\mathbf{u}, p) \qquad \text{Stokes in } \Omega \tag{2.22}$$

$$\mathbf{u} = \mathbf{v}_U \text{ on } U \tag{2.23}$$

$$\mathbf{u} = \mathbf{v}_D \text{ on } D \tag{2.24}$$

$$\mathbf{u}_R - \mathbf{u}_L = \mathbf{g}_u \tag{2.25}$$

$$T(\mathbf{u},p)_R - T(\mathbf{u},p)_L = \mathbf{g}_T. \tag{2.26}$$

By unique continuation from Cauchy data, if $\mathbf{v}_U$ and $\mathbf{v}_D$ are periodic, and we choose $\mathbf{g}_u \equiv \mathbf{0}$ and $\mathbf{g}_T = p_{\text{drive}}\mathbf{n}$, where $\mathbf{n}$ here indicates the normal $(1,0)$ on the L and R walls, then the above BVP is equivalent to (2.14)–(2.18). The special case $\mathbf{v}_U \equiv \mathbf{v}_D \equiv \mathbf{0}$ gives pressure-driven flow in a periodic pipe, free of vesicles. In the general case, there is still a consistency condition on the data. The advantage of the above (non-periodic) BVP in the single unit cell is that the data may be induced by a sum over vesicles which includes only the nearest images and that pressure driving is incorporated naturally.

To solve (2.22)–(2.26), we use a kernel containing only the *near-field* images, plus a small auxiliary basis for smooth Stokes solutions in Ω to account for the effect of the infinite number of *far-field* images. For the latter, we use the "method of fundamental solutions" (MFS) basis [7, 4] of stokeslets with sources lying on a circular contour $\mathcal{C}$ enclosing Ω. (These are also known as "proxy points" [21].) To be precise, the velocity representation is

$$\mathbf{u} \;=\; \mathcal{D}^{\text{near}}_\Gamma \boldsymbol{\tau} + \sum_{m=1}^{M} \mathbf{c}_m \phi_m, \tag{2.27}$$

where

$$(\mathcal{D}^{\text{near}}_\Gamma \boldsymbol{\tau})(\mathbf{x}) := \sum_{|n|\le 1} \int_U D(\mathbf{x}, \mathbf{y}+n\mathbf{d})\boldsymbol{\tau}_U(\mathbf{y})\,ds_\mathbf{y} + \sum_{|n|\le 1} \int_D D(\mathbf{x}, \mathbf{y}+n\mathbf{d})\boldsymbol{\tau}_D(\mathbf{y})\,ds_\mathbf{y} \tag{2.28}$$

is a sum over free-space kernels living on the walls in the central unit cell and its two near neighbors. The second term contains basis functions ϕ_m that satisfy the Stokes equation in the physical domain living in the central unit cell. The basis $\{\phi_m\}$ needs to accurately represent any field due to the "far" periodic copies (i.e. those indexed $\ldots, -3, -2, 2, 3, \ldots$). The source points $\{\mathbf{y}_m\}_{m=1}^M$ are equispaced on a circle of sufficiently large radius R_P centered on the central unit cell, and

$$\phi_m(\mathbf{x}) = S(\mathbf{x}, \mathbf{y}_m), \qquad m = 1, \ldots, M \tag{2.29}$$

is the stokeslet at the pth proxy source. Each coefficient $\mathbf{c}_m \in \mathbb{R}^2$, totalling $2M$ unknowns. This may be viewed as approximating a single-layer density lying on the circle which is able to represent inside any field due to sources lying outside. Since the sources are distant from Ω, the convergence is exponential with a rapid rate; we only need $M = \mathcal{O}(1)$ (typically less than 10^2) *independent of the complexity of the domain or the number of quadrature points needed to accurately represent it.*

The pressure representation corresponding to (2.27) is (summing (2.7) in the same fashion),

$$p = \mathcal{P}_\Gamma^{\text{near}} \boldsymbol{\tau} + \sum_{m=1}^{M} \mathbf{c}_m \varphi_m, \qquad \text{where } \varphi_m(\mathbf{x}) = Q(\mathbf{x}, \mathbf{y}_m), \quad m = 1, \ldots, M. \tag{2.30}$$

Thus, for any density $\boldsymbol{\tau}$ and coefficients $\{\mathbf{c}_m\}$, $(\mathbf{u}, p)$ solves the Stokes equations in Ω.

REMARK 2.2. *There are constraints on the radius R_P: larger R_P allows for more rapid error convergence with respect to M, but if R_P is larger than $3d/2$, then the circle encloses some image sources and the size of the coefficients $\mathbf{c}_m$ grow exponentially large, resulting in catastrophic cancellation. Hence, we fix $R_P = d$ in this study.*

Constructing a linear system is now simply a matter of inserting the representation (2.27) into each of the conditions (2.23)–(2.26), which we now do. Imposing the velocity data on U and D using the jump relation (2.8) (which only affects the $n = 0$ term in (2.28)) gives two coupled boundary integral-algebraic equations,

$$(-\tfrac{1}{2} I + D_{U,U}^{\text{near}}) \boldsymbol{\tau}_U + D_{U,D}^{\text{near}} \boldsymbol{\tau}_D + \sum_{m=1}^{M} \phi_m|_U \mathbf{c}_m = \mathbf{v}_U \qquad \text{on } U \tag{2.31}$$

$$D_{D,U}^{\text{near}} \boldsymbol{\tau}_U + (-\tfrac{1}{2} I + D_{D,D}^{\text{near}}) \boldsymbol{\tau}_D + \sum_{m=1}^{M} \phi_m|_D \mathbf{c}_m = \mathbf{v}_D \qquad \text{on } D\ , \tag{2.32}$$

which we may summarize as

$$A\boldsymbol{\tau} + B\boldsymbol{c} = \boldsymbol{v}\ .$$

Imposing periodicity in matching velocity and traction data (2.25)–(2.26) gives, after noticing cancellations of all of the close wall-wall interactions,

$$(D_{R,U-\mathbf{d}} - D_{L,U+\mathbf{d}}) \boldsymbol{\tau}_U + (D_{R,D-\mathbf{d}} - D_{L,D+\mathbf{d}}) \boldsymbol{\tau}_D + \sum_{m=1}^{M} (\phi_m|_R - \phi_m|_L) \mathbf{c}_m = \mathbf{g}_u \tag{2.33}$$

$$(T_{R,U-\mathbf{d}} - T_{L,U+\mathbf{d}}) \boldsymbol{\tau}_U + (T_{R,D-\mathbf{d}} - T_{L,D+\mathbf{d}}) \boldsymbol{\tau}_D + \sum_{m=1}^{M} \big(T(\phi_m, \varphi_m)|_R - T(\phi_m, \varphi_m)|_L\big) \mathbf{c}_m = \mathbf{g}_T \tag{2.34}$$

which we summarize as

$$C\boldsymbol{\tau} + Q\boldsymbol{c} = \boldsymbol{g}\ .$$

Thus, the four coupled boundary integral-algebraic equations (2.31)–(2.34) may be stacked in pairs and written in a block form

$$\begin{bmatrix} A & B \\ C & Q \end{bmatrix} \begin{bmatrix} \boldsymbol{\tau} \\ \boldsymbol{c} \end{bmatrix} = \begin{bmatrix} \boldsymbol{v} \\ \boldsymbol{g} \end{bmatrix}. \tag{2.35}$$

The roles of the block matrices are as follows: A is a 2nd-kind operator mapping wall densities to (U, D) wall velocities, B maps auxiliary coefficients to wall velocities, C maps wall densities to their discrepancies in the periodicity conditions, and Q maps auxiliary coefficients to their discrepancies. The functions $\boldsymbol{v}$ and $\boldsymbol{g}$ contain the Dirichlet and discrepancy data.

2.3. Discretization of the linear system. The coupled integral-algebraic equations (2.35) need to be discretized: this is performed in a standard fashion using N quadrature nodes $\{\mathbf{x}_{U,j}\}_{j=1}^N$ on U, N nodes $\{\mathbf{x}_{D,j}\}_{j=1}^N$ on D, and K nodes $\{\mathbf{x}_{L,j}\}_{j=1}^K$ on L (the nodes on R being those on L displaced by $\mathbf{d}$). The nodes on U and D are generated using the periodic trapezoid rule applied to a smooth parametrization of the curves, while the L and R wall "collocation" nodes are chosen to be Gauss–Legendre in the vertical coordinate (no weights are needed for these nodes) [35, Ch. 9].

The Nyström method [36, Sec. 12.2] is used to discretize A. For instance, given the quadrature weights $\{w_{U,j}\}_{j=1}^N$ on U, the matrix discretization of the $1,1$ tensor block of $D_{U,U}$ has elements

$$\mathsf{D}_{ij} = \begin{cases} D_{11}(\mathbf{x}_{U,i}, \mathbf{x}_{U,j}) w_{U,j}, & i \neq j \\ -\frac{\kappa(\mathbf{x}_{U,j})}{2\pi} (\mathbf{t}_1(\mathbf{x}_{U,j}))^2 w_{U,j}, & i = j \end{cases} , \tag{2.36}$$

which uses the diagonal limit $\lim_{\mathbf{y}\to\mathbf{x}} D_{ij}(\mathbf{x},\mathbf{y}) = -\frac{\kappa(\mathbf{x})}{2\pi}\mathbf{t}_i(\mathbf{x})\mathbf{t}_j(\mathbf{x})$, where κ is the curvature of the boundary, and $\mathbf{t}$ the unit tangent vector. Other blocks are filled similarly. The result is to replace (2.35) by a discrete linear system of similar structure, which will be solved with a fast direct solver described in Section 4. For more details on a similar periodic discretization scheme, see [12].

3. Application to particulate flows. Now we describe the application of the above periodic BVP solution to vesicle flow simulations, extending the scheme of Veerapaneni et al [53] to periodic flow of vesicle suspensions in rigid pipe-like geometries. We begin with just a single (periodized) vesicle with its boundary γ lying in the periodic unit cell Ω. We solve the steady-state flow problem given forces on the vesicle and then use this solver within the time-stepping scheme.

3.1. Solving for quasi-static fluid flow given the interfacial forces. For simplicity, we consider the case without viscosity contrast. Let $\mathbf{x}(s)$ parametrize the vesicle membrane γ according to arc-length s. The membrane generates forces on the fluid due to bending $\mathbf{f}_b = -\kappa_B \mathbf{x}_{ssss}$ and tension $\mathbf{f}_\sigma = (\sigma \mathbf{x}_s)_s$, where κ_B is the bending modulus and σ is the tension. The total force at each point on γ is then $\mathbf{f} = \mathbf{f}_b + \mathbf{f}_\sigma$. Stress balance and no-slip condition at the membrane-fluid interface imply the jump conditions $[\![T(\mathbf{u},p)]\!]_\gamma = \mathbf{f}$ and $[\![\mathbf{u}]\!]_\gamma = 0$ respectively, where $[\![\cdot]\!]_\gamma$ denotes the jump across γ. In the case of an isolated vesicle in free-space and assuming $\mathbf{f}$ is known, the representation for the fluid velocity $\mathbf{u} = \mathcal{S}_\gamma \mathbf{f}$ and the pressure $p = \mathcal{Q}_\gamma \mathbf{f}$ satisfies these jump conditions as well as the Stokes equations in the bulk.

Our goal in this section is, given only the forces $\mathbf{f}$ on a periodized vesicle $\gamma + n\mathbf{d}$, $n \in \mathbb{Z}$, to solve for the fluid flow $\mathbf{u}$ which results in the periodic confining geometry Ω_Λ, with no-slip boundary conditions and given pressure drop p_{drive} across the channel. The basic idea is to write $\mathbf{u}$ as a sum of the "imposed" flow that the vesicle would generate in an infinite fluid, plus a "response" flow due to the confining geometry Ω. (This is the same concept as the incident and scattered wave in scattering theory [13].) A standard approach in the case $p_{\text{drive}} = 0$ might be to use a periodic imposed flow $\sum_{n\in\mathbb{Z}} \mathcal{S}_{\gamma+n\mathbf{d}}\mathbf{f}$ and for the response to solve the periodic BVP (2.14)–(2.18) with velocity data given by the negative of the imposed flow measured on U and D. The sum of imposed and response flows then would meet our goal. However, this approach has the disadvantage of relying on periodic Greens functions.

Instead, we use the following representation for the physical flow velocity:

$$\mathbf{u} = \mathcal{S}_\gamma^{\text{near}} \mathbf{f} + \mathbf{u}_{\text{resp}}, \qquad \text{where } \mathcal{S}_\gamma^{\text{near}} \mathbf{f} := \sum_{|n|\leq 1} \mathcal{S}_{\gamma+n\mathbf{d}}\mathbf{f} \ . \tag{3.1}$$

The imposed flow (the first term) involves only the vesicle and its immediate neighbor images, as with (2.28). We define the associated imposed pressure similarly: $\mathcal{Q}_\gamma^{\text{near}}\mathbf{f} := \sum_{|n|\leq 1} \mathcal{Q}_{\gamma+n\mathbf{d}}\mathbf{f}$.

The response $\mathbf{u}_{\text{resp}} = \mathbf{u}_{\text{resp}}[\mathbf{f}, p_{\text{drive}}]$ is then the solution to the single-unit-cell BVP (2.22)–(2.26) with the following data involving traces of the imposed flow on the walls:

$$\mathbf{v}_U = -\mathcal{S}_\gamma^{\text{near}}\mathbf{f}|_U \tag{3.2}$$

$$\mathbf{v}_D = -\mathcal{S}_\gamma^{\text{near}}\mathbf{f}|_D \tag{3.3}$$

$$\mathbf{g}_u = -\mathcal{S}_{\gamma-\mathbf{d}}\mathbf{f}|_R + \mathcal{S}_{\gamma+\mathbf{d}}\mathbf{f}|_L \tag{3.4}$$

$$\mathbf{g}_T = -T(\mathcal{S}_{\gamma-\mathbf{d}}\mathbf{f}, \mathcal{Q}_{\gamma-\mathbf{d}}\mathbf{f})|_R + T(\mathcal{S}_{\gamma+\mathbf{d}}\mathbf{f}, \mathcal{Q}_{\gamma+\mathbf{d}}\mathbf{f})|_L + p_{\text{drive}}\mathbf{n} \tag{3.5}$$

It is simple to check that (3.1) then satisfies no-slip velocity data on U and D, is periodic, and the pressure representation has the required pressure drop (2.18). Note that, as in the C block of the previous section, there is cancellation in the discrepancies $\mathbf{g}_u$ and $\mathbf{g}_T$ so that even when vesicles come close to, or intersect, L or R, there are no near-field terms. Effectively, the L and R walls are "invisible" to the vesicles.

To summarize, the algorithm for solving the static periodic pipe flow problem given vesicle forces $\mathbf{f}$ and the driving p_{drive} has three main steps:

i) Evaluate the right-hand side data (3.2)–(3.5), ie

$$\begin{bmatrix} v \\ g \end{bmatrix} = \begin{bmatrix} -S^{\text{near}}_{U,\gamma} \\ -S^{\text{near}}_{D,\gamma} \\ -S_{R,\gamma-\mathbf{d}} + S_{L,\gamma+\mathbf{d}} \\ -K_{R,\gamma-\mathbf{d}} + K_{L,\gamma+\mathbf{d}} \end{bmatrix} \mathbf{f} + \begin{bmatrix} 0 \\ 0 \\ 0 \\ \mathbf{n} \end{bmatrix} p_{\text{drive}} \ ;$$

this will be done with the FMM, except for when the vesicle is close to the wall, in which case a recent spectral close evaluation scheme is used [6].

ii) Solve the rectangular linear system (2.35) for the density τ and coefficient vector c; this is done with the fast direct solver to be described in Section 4.

iii) Evaluate $\mathbf{u}_{\text{resp}}$ (being the solution to (2.22)–(2.26)) using the representation (2.27); this will again be done via the FMM to get $\mathbf{u}_{\text{resp}}|_\gamma$. It is clear that $\mathbf{u}_{\text{resp}}$ is linear both in $\mathbf{f}$ and p_{drive}.

This three-step procedure will become one piece of the following evolution scheme.

3.2. Time-stepping scheme. So far we have only described a quasi-static solution for $\mathbf{u}$ driven by $\mathbf{f}$ and p_{drive}. To close the system, we enforce no-slip conditions on the vesicle, $\dot{\mathbf{x}} = \mathbf{u}|_\gamma$, where $\cdot = \partial/\partial t$. Substituting (3.1) gives the first equation in the integro-differential system of evolution equations for the vesicle dynamics, namely

$$\dot{\mathbf{x}} = S^{\text{near}}_{\gamma,\gamma}\mathbf{f} + \mathbf{u}_{\text{resp}}[\mathbf{f}, p_{\text{drive}}]|_\gamma \tag{3.6}$$

$$0 = \mathbf{x}_s \cdot \dot{\mathbf{x}}_s \tag{3.7}$$

where $\mathbf{f} = -\kappa_B \mathbf{x}_{ssss} + (\sigma \mathbf{x}_s)_s$. Unlike other particulate systems (e.g., drops), the interfacial tension $\sigma(s,t)$ is not known *a priori* and needs to be determined as part of the solution. It serves as a Lagrange multiplier to enforce the local inextensibility constraint—the second equation (3.7) in this system—that the surface divergence of the membrane velocity is zero. This system is driven by p_{drive}, which could vary in time (we take it as constant in our experiments).

The governing equations (3.6)-(3.7) are numerically stiff owing to the presence of high-order spatial derivatives in the bending force. As shown in [53], explicit time-stepping schemes, such as the forward Euler method, suffer from a third-order constraint on the time step size, rendering them prohibitively expensive for simulating vesicle suspensions. Therefore, we use the semi-implicit scheme formulated in [53] with a few modifications to improve the overall numerical accuracy and stability. Given a time step size Δt and the membrane position and tension at the kth time step, $(\mathbf{x}^k, \sigma^k)$, we evolve to $(\mathbf{x}^{k+1}, \sigma^{k+1})$ by using a first-order semi-implicit time-stepping scheme on (3.6)-(3.7). For implementational convenience, however, we treat the discretized membrane velocity, denoted with a slight abuse of notation by $\mathbf{u} = (\mathbf{x}^{k+1} - \mathbf{x}^k)/\Delta t$, as the unknown instead of $\mathbf{x}^{k+1}$. The scheme, then, is given by

$$\mathbf{u} - S^{\text{near}}_{\gamma,\gamma}[-\Delta t \kappa_B \mathbf{u}_{ssss} + (\sigma^{k+1}\mathbf{x}^k_s)_s] = S^{\text{near}}_{\gamma,\gamma}[-\kappa_B \mathbf{x}^k_{ssss}] + \mathbf{u}_{\text{resp}}[-\kappa_B \mathbf{x}^k_{ssss} + (\sigma^k \mathbf{x}^k_s)_s, p_{\text{drive}}]|_\gamma \tag{3.8}$$

$$\mathbf{x}^k_s \cdot \mathbf{u}_s = 0. \tag{3.9}$$

Since the bending force is a nonlinear function of the membrane position, the standard principle of semi-implicit schemes—to treat the terms with highest-order spatial derivatives implicitly [3]—has been applied to the particular linearization. The tension is treated implicitly and the vesicle-channel interaction, explicitly[1]. The single-layer operator $S^{\text{near}}_{\gamma,\gamma}[\cdot]$ as well as the differential operator $(\cdot)_s$ are constructed using

[1]When the vesicle is located very close to the channel, say $\mathcal{O}(h)$ away where h is the lowest distance between spatial grid points on the vesicle, a semi-implicit treatment of $\mathbf{u}_{\text{resp}}$ would be more efficient since the interaction force also induces numerical stiffness in this scenario. Such a scheme, however, requires non-trivial modifications to our fast direct solver of Section 4. Therefore, we postponed this exercise to future work.

$\mathbf{x}^k$. In summary, we solve the following linear system for the unknowns $(\mathbf{u}, \sigma^{k+1})$:

$$\begin{bmatrix} I + \Delta t \kappa_B S^{\text{near}}_{\gamma,\gamma} \partial_{ssss} & -S^{\text{near}}_{\gamma,\gamma} \partial_s(\mathbf{x}^k_s \cdot) \\ \mathbf{x}^k_s \cdot \partial_s & 0 \end{bmatrix} \begin{bmatrix} \mathbf{u} \\ \sigma^{k+1} \end{bmatrix} = \begin{bmatrix} -\kappa_B S^{\text{near}}_{\gamma,\gamma} \mathbf{x}^k_{ssss} + \mathbf{u}_{\text{resp}}|_\gamma \\ 0 \end{bmatrix} \tag{3.10}$$

with given $(\mathbf{x}^k, \sigma^k)$ and then update the membrane positions as $\mathbf{x}^{k+1} = \mathbf{x}^k + \Delta t \mathbf{u}$. The operators are discretized using a spectrally-accurate Nyström method (with periodic Kress corrections for the log singularity [36, Sec. 12.3]) for the single-layer operator and a standard periodic spectral scheme for the differentiation operators. The resulting discrete linear system is solved via GMRES, with all distant interactions applied using the Stokes FMM (e.g., see Appendix D of [53]).

On the initial time step, we generally set $\sigma_0 = 0$. One could obtain an improved initial guess for σ_0 by setting $\sigma_0 = 0$ and solving (3.8) and (3.9) for σ with $\Delta t = 0$. This could be thought of as the first iteration of a Picard iteration for σ. To enable long-time simulations, we incorporate three supplementary steps in our time-stepping scheme. First, we modify the constraint equation (3.9) as[2]

$$\mathbf{x}^k_s \cdot \mathbf{u}_s = \frac{L_0 - L_k}{\Delta t L_k}, \tag{3.11}$$

where L_k represents the perimeter of the vesicle at the kth time step. We will refer to this as the "arc length correction (ALC)." We prove in Appendix A that without the ALC, the perimeter of a vesicle will increase monotonically with the number of time steps (total error still scales as $\mathcal{O}(\Delta t)$). This would mean that the vesicle's reduced area can become very low when a large number of time steps are taken. Consequently, the simulated dynamics may correspond to a totally different system than what was originally intended (e.g., a tank-treading vesicle in shear flow might tumble if the reduced area is lowered enough). Executing the ALC at every time step, on the other hand, guarantees a $\mathcal{O}(\Delta t^2)$ convergence rate, but more importantly, the error is independent of the number of time steps for a fixed Δt (see Theorem A.2). Second, we correct the error incurred in the enclosed area of the vesicle after every time step by solving a quadratic equation in one variable (see Appendix B). Finally, we reparameterize γ at every time step so that spatial discretization points are located *approximately* equal arc lengths apart (see Appendix C).

Suspension flow. Although we have presented the case with a single (periodized) vesicle γ, the above scheme carries over naturally to multiple vesicles. Since such extension has been described previously in other contexts (e.g., see [53] for free-space and [48] for constrained geometry problems) and does not modify in any way our periodization scheme, we only highlight the main steps here. First, the single-layer potential in the representation (3.1) is replaced with a sum of such potentials over all of the vesicles. The equations for the imposed flow data on the walls (3.2)–(3.5) are then modified accordingly. In discretizing the evolution equation for each individual vesicle, the bending force in the self-interaction term is treated semi-implicitly, similar to (3.8)–(3.9). However, the bending forces in the vesicle-vesicle interaction terms can either be treated semi-implicitly, resulting in a dense linear system, or explicitly, resulting in a block tri-diagonal system. While the latter scheme has a marginally smaller computational cost, the former scheme has better stability properties in general since vesicle-vesicle interactions can induce stiffness into the evolution equations when they are located close to each other. In our implementation, we treat all vesicle interactions semi-implicitly. We use a recent single-layer close evaluation scheme [6] to compute the nearby vesicle interactions, whereas for distant interactions, we use the FMM.

4. Fast direct solver for the fixed channel geometry. At every time step, the right-hand side in (3.10) must be evaluated, which involves the channel response 3-step solution given at the end of Section 3.1. However, since the channel geometry is fixed, a fast direct solver enables the second, potentially most expensive, of these three steps to be performed in $\mathcal{O}(N)$ time with a very small constant. Recall that this step involves solving the 2×2 block integral equation system (2.35). Upon discretization, this becomes a rectangular system of size $2(N+K) \times 2(N+M)$ given by

$$\begin{bmatrix} \mathsf{A} & \mathsf{B} \\ \mathsf{C} & \mathsf{Q} \end{bmatrix} \begin{bmatrix} \boldsymbol{\tau} \\ \boldsymbol{c} \end{bmatrix} = \begin{bmatrix} \boldsymbol{v} \\ \boldsymbol{g} \end{bmatrix}. \tag{4.1}$$

[2]The main idea here is similar in spirit to the correction formula applied in [52], but in our scheme, we do not introduce any penalty parameters and also rigorously prove its convergence rate (Appendix A).

This section presents a fast direct solution technique for (4.1). The idea is to precompute for $\mathcal{O}(N)$ computational cost, the factors in the block matrix 2×2 solve. Then, the contribution from the channel geometry in the time-stepping scheme only requires a collection of inexpensive linear-scaling matrix vector multiplies.

This section begins by presenting the 2×2 block solve. The remainder of the section describes how to efficiently build and apply the block solver. The bulk of the novelty in this work lies in the linear scaling technique for representating the interactions of neighboring geometries.

4.1. The block solve. The solution to the rectangular system (4.1) is given by

$$\begin{aligned} \boldsymbol{c} &= -\mathsf{S}^{\dagger}\left(\boldsymbol{g} - \mathsf{C}\mathsf{A}^{-1}\boldsymbol{v}\right) \\ \boldsymbol{\tau} &= \mathsf{A}^{-1}\boldsymbol{v} - \mathsf{A}^{-1}\mathsf{B}\boldsymbol{c}, \end{aligned}$$

where $\mathsf{S} = \mathsf{Q} - \mathsf{C}\mathsf{A}^{-1}\mathsf{B}$, and $\mathsf{S}^{\dagger}$ denotes the pseudoinverse of S. S is often referred to as a Schur complement matrix.

The Schur complement and $\mathsf{A}^{-1}\mathsf{B}$ need only be formed once, independent of the number of right-hand sides (i.e. time steps). Beyond that, the array $\mathsf{A}^{-1}\boldsymbol{v}$ need only be computed once per solve.

When a large number of points N are needed to discretize the walls due to a complex geometry or small vesicle size, the cost of computing the block solve is dominated by the cost of computing the inverse of A. When N is large, computing A^{-1} is computationally prohibitive. Fortunately, the matrix A has structure which can be exploited to reduce the computational cost of the block solve.

Recall that A is the discretization of (2.28) added to $-\frac{1}{2}\mathsf{I}$. The underlying structure of A can exploited first by considering its expanded form given by

$$\mathsf{A} = \mathsf{A}_0 + (\mathsf{A}_{-1} + \mathsf{A}_1), \tag{4.2}$$

where A_j corresponds to the discretization of the self ($j = 0$) and neighbor ($j = -1, 1$) channel geometry interactions.

Since the majority of the discretization points on the neighboring geometries are well-separated, potential theory states that A_{-1} and A_1 are low rank (i.e. A_j has a rank l, where $l \ll N$). Thus, each matrix admits a factorization of the form $\mathsf{A}_j = \mathsf{L}_j\mathsf{R}_j$, where L_j and R_j are of size $N \times l$ for $j = -1, 1$. Thus,

$$\mathsf{A} = \mathsf{A}_0 + \mathsf{L}\mathsf{R}$$

, where $\mathsf{L} = [\mathsf{L}_{-1}|\mathsf{L}_1]$ and $\mathsf{R} = \left[\mathsf{R}_{-1}^T|\mathsf{R}_1^T\right]^T$.

A consequence of utilizing the low rank factorization is that the inverse of A can be computed via the *Sherman-Morrison-Woodbury formula*

$$(\mathsf{A}_0 + \mathsf{L}\mathsf{R})^{-1} = \mathsf{A}_0^{-1} - \mathsf{A}_0^{-1}\mathsf{L}\left(\mathsf{I} + \mathsf{R}\mathsf{A}_0^{-1}\mathsf{L}\right)^{-1}\mathsf{R}\mathsf{A}_0^{-1}. \tag{4.3}$$

Note, only the inverse of A_0 and $\left(\mathsf{I} + \mathsf{R}\mathsf{A}_0^{-1}\mathsf{L}\right)$ need to be computed. Section 4.3 describes a technique for inverting A_0 with a cost that scales linearly with N for most wall geometries. Other fast inversion techniques, such as [57, 50, 56, 8, 9, 21, 2], can be utilized in place of the method in Section 4.3. The matrix $\left(\mathsf{I} + \mathsf{R}\mathsf{A}_0^{-1}\mathsf{L}\right)$ is $2l \times 2l$ in size and is small enough to be inverted rapidly via dense linear algebra.

4.2. Construction of low-rank factorization for neighbor interactions. The cost of constructing the factorizations of A_{-1} and A_1 using general linear algebra techniques, such as QR, is $\mathcal{O}(N^2 l)$ and thus would negate the key reduction in asymptotic complexity gained by using fast direct solvers, such as the one described in Section 4.3 to invert A_0. To maintain the optimal asymptotic complexity, we utilize ideas from potential theory, similar to those in [20]. Unlike [20], the neighboring geometries touch the geometry in the unit cell. As a result, a new technique for computing the factorizations is required. For simplicity of presentation, we describe the new low-rank factorization technique for the

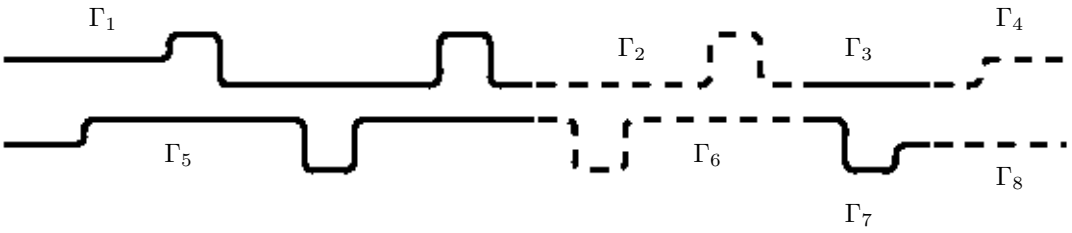

FIG. 4.1. *Illustration of the dyadic partitioning used to compress the interaction of the walls in the unit cell with the right neighbor, i.e.* A_1.

matrix $\mathsf{A}_1 = \mathsf{L}_1\mathsf{R}_1$. The technique is applied in a similar manner to construct the low-rank factorization of $\mathsf{A}_{-1} = \mathsf{L}_{-1}\mathsf{R}_{-1}$. As in [20], an interpolatory decomposition [24, 11] is utilized.

DEFINITION 4.1. *The interpolatory decomposition of an* $m \times n$ *matrix* M *that has rank* l *is the factorization*

$$\mathsf{M} = \mathsf{P}\mathsf{M}(J(1:l),:),$$

where J *is a vector of integers* j_i *with* $1 \leq j_i \leq m$*, and* P *is an* $m \times l$ *matrix that contains an* $l \times l$ *identity matrix. Namely,* $\mathsf{P}(J(1:l),:) = \mathsf{I}_l$.

First, the upper and lower part of the geometry is partitioned into a collection of M sections Γ_j via dyadic refinement where the rectangular boxes enclosing each segment get smaller as they approach the neighboring cells. Thus, $\partial\Omega_0 = \cup_{j=1}^{M}\Gamma_j$, where M is the number of sections $\partial\Omega_0$ is partitioned into. Figure 4.1 illustrates the partitioning of the walls for the channel with reservoirs in Figure 5.1(d) when compressing the interaction with the right neighbor Ω_1. The refinement is stopped when the smallest box contains no more than a specified number of points $n_{\max}$. Typically, $n_{\max} = 45$ is a good choice.

For each portion of the boundary Γ_j that is not touching Ω_1, consider a circle concentric with a box bounding Γ_j with radius slightly less than the distance from the center of the box enclosing Γ_j to the "wall" where Ω_0 and Ω_1 meet. From potential theory, we know that any field generated by sources outside of this circle can be approximated well by placing enough equivalent charges on the circle. In practice, it is possible to place a small number of "proxy" points spaced evenly on the circle. For the experiments in this paper, we found it is sufficient to have 80 proxy points. Figure 4.2 (a) and (b) illustrate the proxy points for Γ_1 and Γ_6. An interpolatory decomposition is then found for the matrix denoted A^{proxy} that characterizes the interactions between Γ_j and the proxy points. The result is a matrix P_j and index vector J_j. For Γ_j touching $\partial\Omega_1$, a collection of 80 proxy points are placed evenly on a circle with a radius 1.75 times the radius of the smallest circle enclosing Γ_j. All of the points on $\partial\Omega_1$ that lie inside this proxy circle are called *near* points. Figures 4.2(c) and (d) illustrates the proxy and near points for Γ_4. An interpolatory decomposition is then formed for the matrix

$$[\mathsf{A}_1(\Gamma_j, I^{\text{near}})|A^{\text{proxy}}]$$

, where I^{near} corresponds to the indices of the portion of $\partial\Omega_1$ that are near Γ_j. The points $\partial\Omega_0$ picked by the interpolatory decomposition are called *skeleton points.*

We can now assemble the factors $\mathsf{L}_1\mathsf{R}_1$ by sweeping through the regions Γ_j. Let $J = [J_1(1 : l_1), \ldots, J_M(1 : l_M)]$. Then the matrix L_1 is a block diagonal matrix where each block is a P_j and $\mathsf{R}_1 = \mathsf{A}_1(J,:)$.

REMARK 4.1. *The factorization as described does not have optimal rank. Optimal rank can be obtained by recompressing via additional low rank factorizations while assembling* L_1 *and* R_1*. Depending on the geometry and the computer, it may or may not be beneficial to do the recompression.*

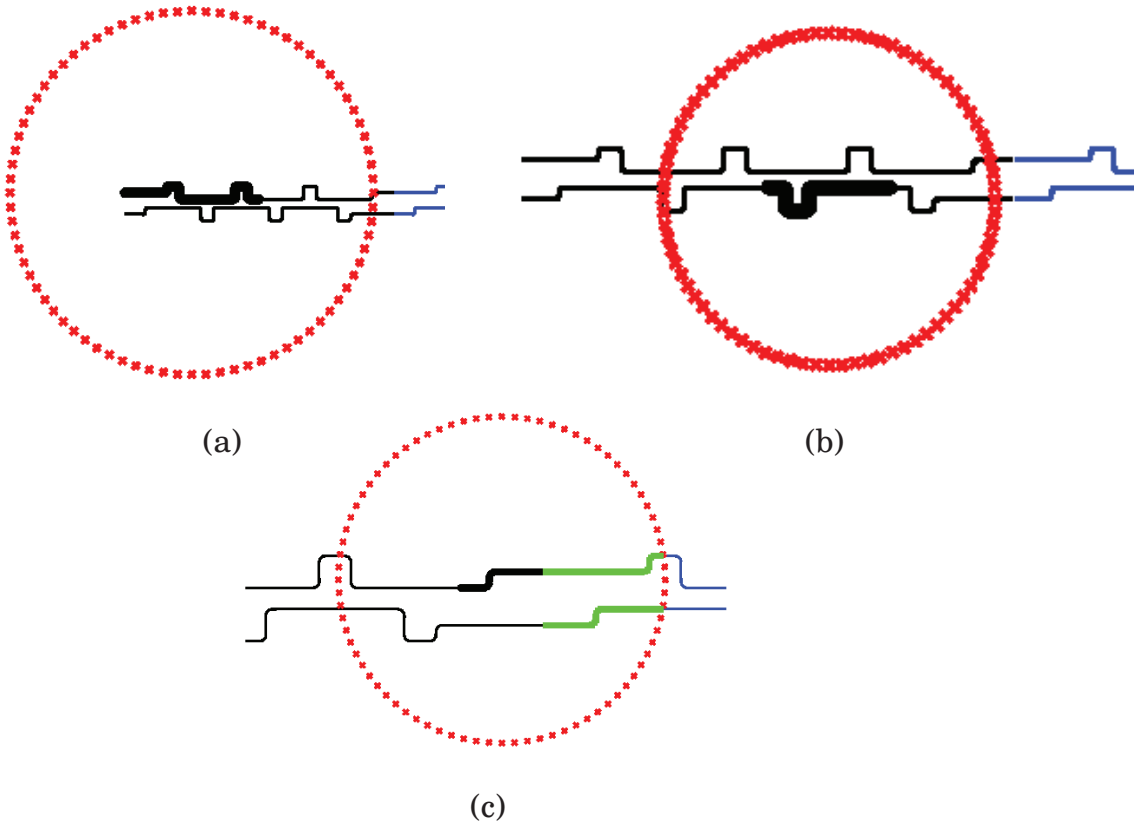

FIG. 4.2. *(a) and (b) illustration of the proxy points ($\boldsymbol{x}$) for Γ_1, Γ_6. (c) illustrates the proxy points and the near points (thick green line) on $\partial\Omega_1$ for Γ_4.*

4.3. HBS inversion. As stated previously, the dense matrix A_0 has structure that we call *Hierarchically Block Separable (HBS)*. The HBS structure allows for an approximation of A_0^{-1} to be computed rapidly. Loosely speaking, its off-diagonal blocks are low rank. This arises because A_0 is the discretization on a curve of an integral operator with smooth kernel (when the walls are not space-filling). This section briefly describes the HBS property and how it can be exploited to rapidly construct an approximate inverse of a matrix. For additional details, see [21]. Note that the HBS property is very similar to the concept of *Hierarchically Semi-Separable (HSS)* matrices [50, 10].

4.3.1. Block separable. Let M be an $mp \times mp$ matrix that is blocked into $p \times p$ blocks, each of size $m \times m$.

We say that M is "block separable" with "block-rank" k if for $\tau = 1, 2, \ldots, p$, there exist $n \times k$ matrices U_τ and V_τ such that each off-diagonal block $\mathsf{M}_{\sigma,\tau}$ of M admits the factorization

$$\underset{m\times m}{\mathsf{M}_{\sigma,\tau}} = \underset{m\times k}{\mathsf{U}_\sigma}\ \underset{k\times k}{\tilde{\mathsf{M}}_{\sigma,\tau}}\ \underset{k\times m}{\mathsf{V}_\tau^*}, \qquad \sigma,\tau \in \{1, 2, \ldots, p\}, \quad \sigma \neq \tau. \tag{4.4}$$

Observe that the columns of U_σ must form a basis for the columns of all off-diagonal blocks in row σ, and analogously, the columns of V_τ must form a basis for the rows in all of the off-diagonal blocks in column τ. When (4.4) holds, the matrix M admits a block factorization

$$\underset{mp\times mp}{\mathsf{M}} = \underset{mp\times kp}{\mathsf{U}}\ \underset{kp\times kp}{\tilde{\mathsf{M}}}\ \underset{kp\times mp}{\mathsf{V}^*} + \underset{mp\times mp}{\mathsf{D}}, \tag{4.5}$$

where

$$\mathsf{U} = \operatorname{diag}(\mathsf{U}_1, \mathsf{U}_2, \ldots, \mathsf{U}_p), \quad \mathsf{V} = \operatorname{diag}(\mathsf{V}_1, \mathsf{V}_2, \ldots, \mathsf{V}_p), \quad \mathsf{D} = \operatorname{diag}(\mathsf{D}_1, \mathsf{D}_2, \ldots, \mathsf{D}_p),$$

and

$$\tilde{\mathsf{M}} = \begin{bmatrix} 0 & \tilde{\mathsf{M}}_{12} & \tilde{\mathsf{M}}_{13} & \cdots \\ \tilde{\mathsf{M}}_{21} & 0 & \tilde{\mathsf{M}}_{23} & \cdots \\ \tilde{\mathsf{M}}_{31} & \tilde{\mathsf{M}}_{32} & 0 & \cdots \\ \vdots & \vdots & \vdots & \end{bmatrix}.$$

Once the matrix M has been put into block separable form, its inverse is given by

$$\mathsf{M}^{-1} = \mathsf{E}\,(\tilde{\mathsf{M}} + \hat{\mathsf{D}})^{-1}\,\mathsf{F}^* + \mathsf{G}, \tag{4.6}$$

where

$$\hat{\mathsf{D}} = \left(\mathsf{V}^*\,\mathsf{D}^{-1}\,\mathsf{U}\right)^{-1}, \tag{4.7}$$

$$\mathsf{E} = \mathsf{D}^{-1}\,\mathsf{U}\,\hat{\mathsf{D}}, \tag{4.8}$$

$$\mathsf{F} = (\hat{\mathsf{D}}\,\mathsf{V}^*\,\mathsf{D}^{-1})^*, \tag{4.9}$$

$$\mathsf{G} = \mathsf{D}^{-1} - \mathsf{D}^{-1}\,\mathsf{U}\,\hat{\mathsf{D}}\,\mathsf{V}^*\,\mathsf{D}^{-1}. \tag{4.10}$$

4.3.2. Hierarchically Block-Separable. Informally speaking, a matrix M is *Hierarchically Block-Separable* (HBS) if it is amenable to a *telescoping* version of the above block factorization. In other words, in addition to the matrix M being block separable, so is $\tilde{\mathsf{M}}$ once it has been re-blocked to form a matrix with $p/2 \times p/2$ blocks, and one is able to repeat the process in this fashion multiple times.

For example, a "3 level" factorization of M is

$$\mathsf{M} = \mathsf{U}^{(3)}\big(\mathsf{U}^{(2)}\big(\mathsf{U}^{(1)}\,\tilde{\mathsf{M}}^{(0)}\,(\mathsf{V}^{(1)})^* + \mathsf{B}^{(1)}\big)(\mathsf{V}^{(2)})^* + \mathsf{B}^{(2)}\big)(\mathsf{V}^{(3)})^* + \mathsf{D}^{(3)}, \tag{4.11}$$

where the superscript denotes the level.

The HBS representation of an $N \times N$ matrix requires $\mathcal{O}(Nk)$ to store and to apply to a vector. By recursively applying formula (4.6) to the telescoping factorization, an approximation of the inverse can be computed with $\mathcal{O}(Nk^2)$ computational cost; see [21]. This compressed inverse can be applied to a vector (or a matrix) very rapidly.

5. Numerical results. In this section, we test the performance of our algorithm on simulating flows through four geometries ranging in complexity from a flat channel to a more complicated space-filling geometry. We apply a pressure-drop on each side which drives the flow in the positive x-direction while simultaneously enforcing no-slip boundary conditions on the walls. The streamlines are plotted in Fig. 5.1 along with the magnitude of the velocity which is indicated by the color in the background. We will refer to geometry (a) as the straight channel, (b) as the converging-diverging channel, (c) as the serpentine channel, and (d) as the channel with reservoirs. Note that all the wall geometries are $\mathcal{C}^\infty$ curves, constructed using partitions of unity to avoid sharp corners. We used the trapezoidal rule, with $N = 4{,}000$ nodes on each wall, to evaluate the layer potentials at interior targets for this test[3]

In the following subsection, we analyze the convergence of the periodization scheme as we vary the number of proxy points (denoted by M), discretization points on L and R (denoted by K), and quadrature points on the walls (denoted by N). For these tests, we set the boundary conditions and pressure-drop such that the horizontal Poiseuille flow $\mathbf{u}_e(\mathbf{x}) = (\alpha x_2^2, 0)$, $p_e(\mathbf{x}) = \alpha\mu x_1$, with $\alpha = p_{\text{drive}}/(2\mu d)$ satisfies the BVP. For all tests, we use $\alpha = 0.2$, $\mu = 0.7$, and $d = 2\pi$. The relative errors are defined as

$$\text{relative error in velocity} = \frac{\|\mathbf{u} - \mathbf{u}_e\|_2}{\max\limits_{\mathbf{x}\in\Omega}\left(\|\mathbf{u}_e\|_2\right)} \tag{5.7}$$

$$\text{relative error in pressure} = \left|\frac{p - p_e}{p_{\text{drive}}}\right|, \tag{5.8}$$

[3]For some applications, the trapezoidal rule may not be an ideal choice since it cannot yield uniform convergence for points that are very close to the walls. For example, notice the thin band near the wall in Fig. 5.4. Accurate close-evaluation schemes must be used instead. In our setting, the no-slip condition at the walls leads to low velocities near the walls and in general helps alleviate numerical instabilities naturally. Incorporating close evaluation schemes (for walls) is one of the first tasks in our future work.

Algorithm 1 Main (single vesicle)

Step 1: Compress A^{-1} using the fast direct solver
for $k = 0 : N_{\Delta t} - 1$ **do**
Step 2: Compute bending and tension forces ($\mathbf{f}_\sigma = \mathbf{0}$ if $k = 0$)

$$\mathbf{f}_b = -\kappa_B \mathbf{x}^k_{ssss}, \quad \mathbf{f}_\sigma = \left(\sigma^k \mathbf{x}^k_s\right)_s, \quad \mathbf{f} = \mathbf{f}_b + \mathbf{f}_\sigma \tag{5.1}$$

Step 3: Compute vesicle-wall and vesicle-side interactions (step (i) from Section 3.1)

$$\mathbf{v}_U = -\mathcal{S}^{\text{near}}_\gamma \mathbf{f}\big|_U \tag{5.2}$$
$$\mathbf{v}_D = -\mathcal{S}^{\text{near}}_\gamma \mathbf{f}\big|_D \tag{5.3}$$
$$\mathbf{g}_u = -\mathcal{S}_{\gamma - \mathbf{d}} \mathbf{f}\big|_R + \mathcal{S}_{\gamma + \mathbf{d}} \mathbf{f}\big|_L \tag{5.4}$$
$$\mathbf{g}_T = -T\left(\mathcal{S}_{\gamma - \mathbf{d}} \mathbf{f}, \mathcal{Q}_{\gamma - \mathbf{d}} \mathbf{f}\right)\big|_R + T\left(\mathcal{S}_{\gamma + \mathbf{d}} \mathbf{f}, \mathcal{Q}_{\gamma + \mathbf{d}} \mathbf{f}\right)\big|_L + p_{\text{drive}} \mathbf{n} \tag{5.5}$$

Step 4: Solve for $\boldsymbol{\tau}$, $\mathbf{c}$ using the fast direct solver (step (ii) from Section 3.1)
Step 5: Compute wall-vesicle and proxy-vesicle interactions (step (iii) from Section 3.1)

$$\mathbf{u}_{\text{resp}} = \mathcal{D}^{\text{near}}_\Gamma \boldsymbol{\tau}\big|_\gamma + \sum_{m=1}^{M} \mathbf{c}_m \phi_m\big|_\gamma \tag{5.6}$$

Step 6: Solve (3.10) for $\mathbf{u}$ and σ^{k+1} using GMRES with the constraint equation modified as (3.11).
Step 7: Set $\mathbf{x}^{k+1} = \mathbf{x}^k + \Delta t\, \mathbf{u}$
Step 8: Apply area correction to $\mathbf{x}^{k+1}$ (Appendix B)
Step 9: Apply the reparameterization scheme on $\mathbf{x}^{k+1}$ (Algorithm 2 of Appendix C)
end for

where $(\mathbf{u}, p)$ is the numerical solution obtained by the algorithm. We report the maximum relative error at the same three points $(0.8, 0)$, $(3.2, 0)$, and $(4.8, 0)$ for all four geometries (not too close to the walls).

5.1. Periodization scheme. Here, we analyze the convergence properties of the periodization scheme. We begin by focusing on the sides L and R and the proxy sources $\mathcal{C}$. We compute the relative errors for the velocity in the converging-diverging channel as the number of side points K and proxy sources M are varied. The number of quadrature points along the walls is fixed to $N = 600$, which, based on previous experimentation, is enough to resolve the flow to near machine precision. As can be seen in Fig. 5.2, rapid convergence is achieved and high accuracy is obtained using a relatively small K and M. This is mainly because of the analytical cancellation of close interactions of L, R, and $\mathcal{C}$ with the walls, as shown in equations (2.33) and (2.34). Since flow far from a disturbance is generally smooth with rapidly decaying Fourier modes, the distant interactions relating to L, R, and $\mathcal{C}$ are easily resolved.

Next, we investigate how the complexity of the geometry affects the convergence rate of the side points and the proxy sources. This is done by setting $M = 4K$ and measuring the relative errors pertaining to the four test geometries as K is varied. Fig. 5.3(a) shows the results. Notice that once we account for the height of the inlet/outlet, the convergence of all four geometries is almost exactly the same, independent of the complexity of the channel. This feature is especially nice for problems dealing with multi-scale physics, such as our application to vesicle flows.

Finally, we focus on the walls U and D. Fig. 5.3(b) shows the relative errors for all four test geometries as the number of quadrature points N is varied. In all cases, we observe spectral convergence. (Recall that the geometries from Fig. 5.1 were constructed using partitions of unity and are C^∞.) The relative errors for both the velocity and pressure inside the serpentine channel are plotted in Fig. 5.4. The small band of errors near the wall is the result of using the trapezoidal rule as opposed to a close evaluation scheme for the walls. It will turn out that in some cases this error band may cause problems for the vesicle simulation if vesicles drift too close to the walls, as we will discuss in the following section.

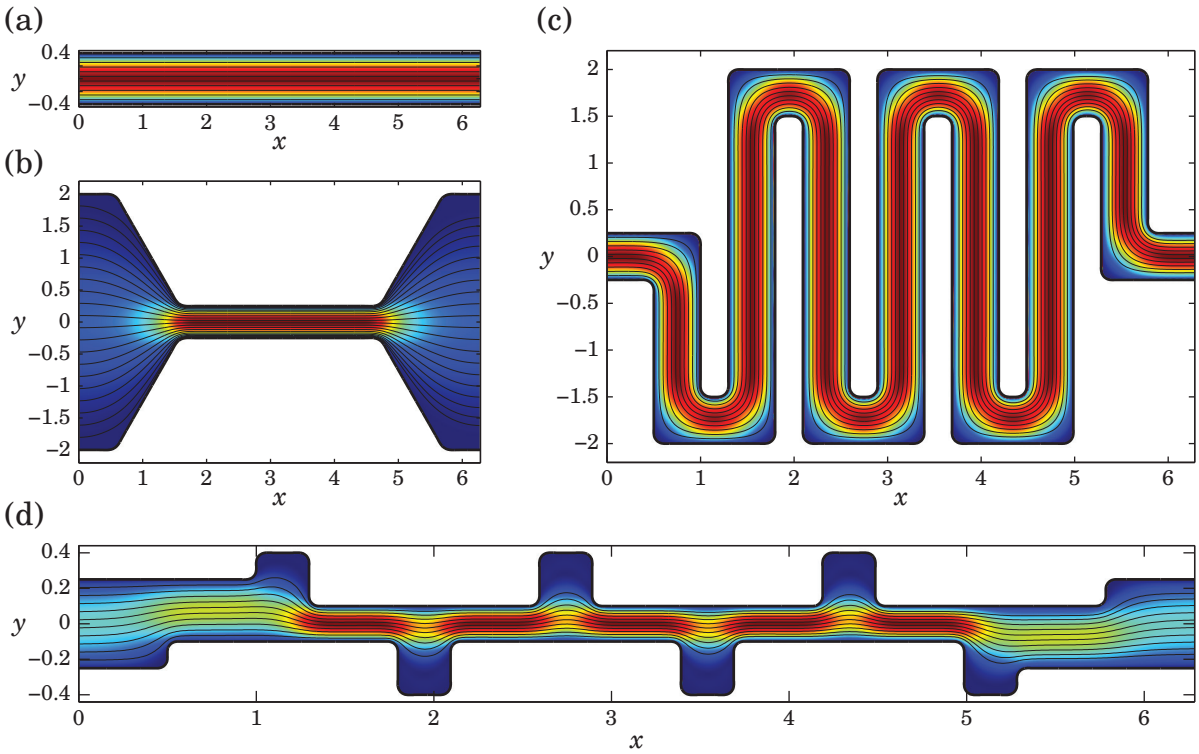

FIG. 5.1. *Streamlines for pressure-driven flows through four geometries with no-slip boundary conditions. The background color indicates the magnitude of the velocity (red indicates high and blue indicates low). For the remainder of this paper, we will refer to geometry (a) as the straight channel, (b) as the converging-diverging channel, (c) as the serpentine channel, and (d) as the channel with reservoirs. In all four geometries, we use* $N = 4{,}000$ *quadrature points on each wall,* $K = 32$ *points on each side, and* $M = 128$ *proxy points. The ring of proxy sources has radius* $d = 2\pi$.

5.2. Vesicle flow simulation. We now give numerical results for the vesicle flow simulation described in Section 3. A snapshot of 109 vesicles flowing through the serpentine channel is shown in Fig. 5.5(a), where a pressure difference between the left and right sides drives the flow in the positive x-direction. When modeling such flows, it's vitally important that vesicles avoid collisions with other vesicles and the walls. We have found that using 64 points per vesicle along with a close evaluation scheme is usually enough to prevent vesicle-vesicle collisions as long as vesicle shapes are not elongated and the time step is not too big. For the simulations in Figs. 1.1(a) and 5.5(a), the time step was set to

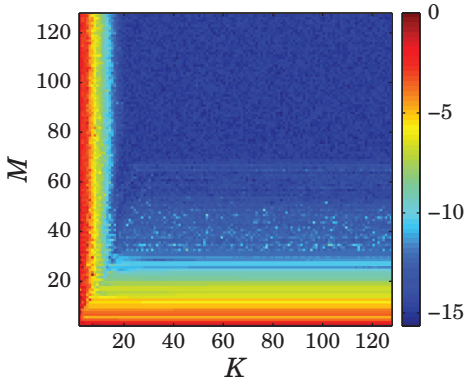

FIG. 5.2. *Logarithm of the relative errors for the velocity as we vary the number of points on the sides and on the ring of proxy sources for the converging-diverging geometry in Fig. 5.1(b). For each test point, the velocity was computed using* $N = 600$ *quadrature points per wall. The boundary integral equation was solved using GMRES.*

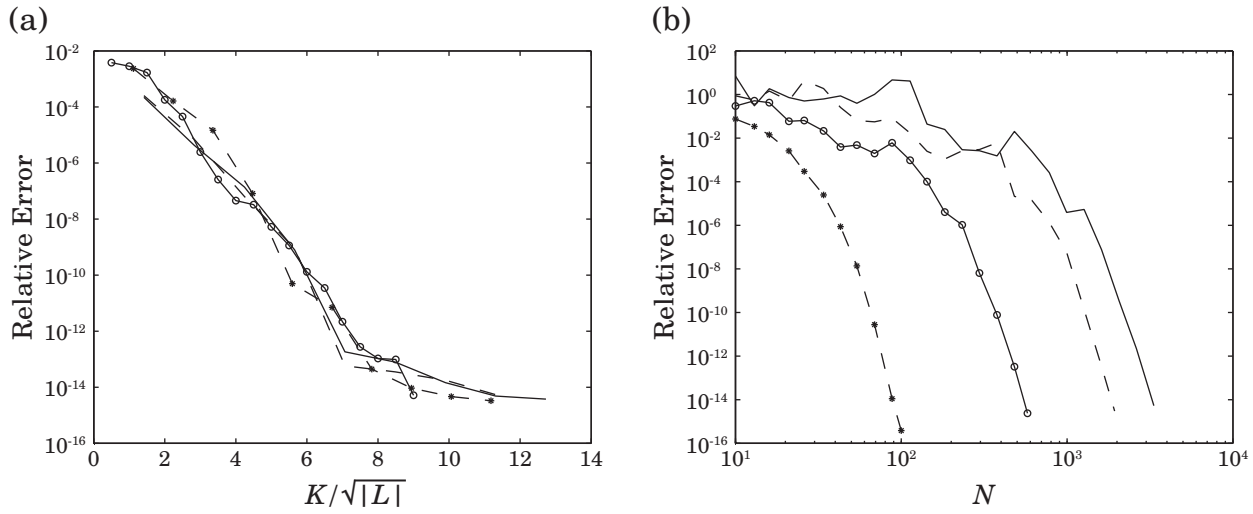

FIG. 5.3. (∗) *straight channel,* (○) *converging-diverging channel,* (−) *serpentine channel,* (−−) *channel with reservoirs. (a) Relative errors as the number of side points and proxy sources are varied. For each geometry, we set the number of points on the walls to* $N = 4{,}000$ *and the number of proxy sources to* $M = 4K$*. The height of the inlet/outlet is denoted by* $|L|$*. Notice that the convergence appears to be independent of the complexity of the channel. (b) Relative errors as the number of points on the walls are varied. In each case, we used* $K = 32$ *points per side and* $M = 128$ *proxy sources. Spectral convergence was obtained for all four geometries.*

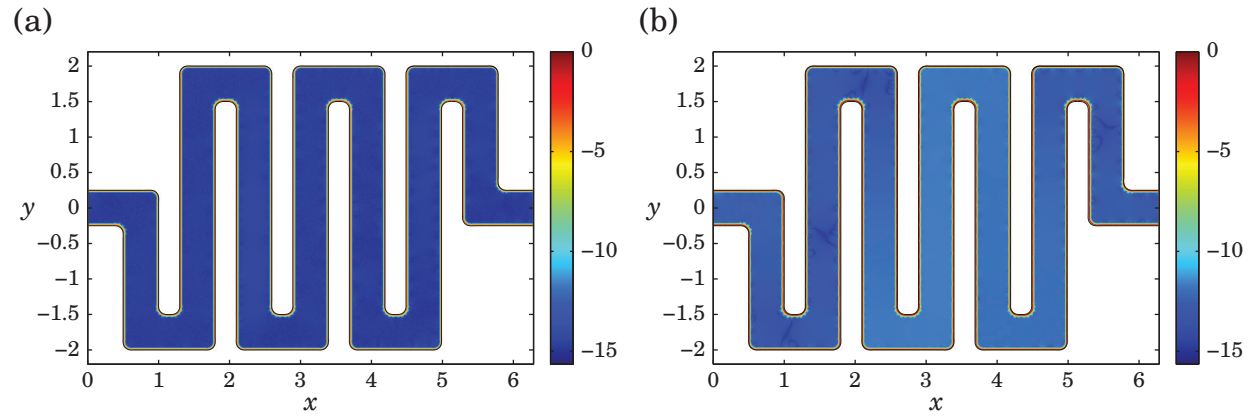

FIG. 5.4. *Spatial plot of the logarithm of relative errors in velocity (a) and pressure (b). Errors were computed using the trapezoidal rule with 4,000 quadrature points per wall. The thin band near the walls, where low accuracy is obtained, may be resolved using a close evaluation scheme.*

$\Delta t = 0.005$. When modeling pressure driven flows with no-slip boundary conditions, the vesicles often keep a safe distance from the walls and a close evaluation scheme is not always required. However, there are plenty of situations where this is not the case. In general, vesicle-wall collisions tend to occur more often when the channel geometry has bends or tight spaces, the number of vesicles is large, or the bending modulus κ_B is high.

We now focus on the performance of the algorithm. Fig. 5.5(b) shows the scaling as the number of vesicles in the serpentine geometry is increased. When performing the timings, the vesicle dimensions were scaled to maintain the same relative spacing. Each vesicle had 64 discretization points and the number of quadrature points on the walls was set to 29 times the number of vesicles. The algorithm maintained linear scaling to over 1,000 vesicles on a laptop with a 2.4 GHz dual-core Intel Core i5 processor and 8 GB of RAM. In Table 5.1, we give the CPU time distribution for the simulation with 1,020 vesicles. Approximately 82% of the computational time was spent computing vesicle-vesicle, vesicle-wall, or wall-vesicle interactions. The majority of this expense was handled by the FMM. Approximately 7% of the

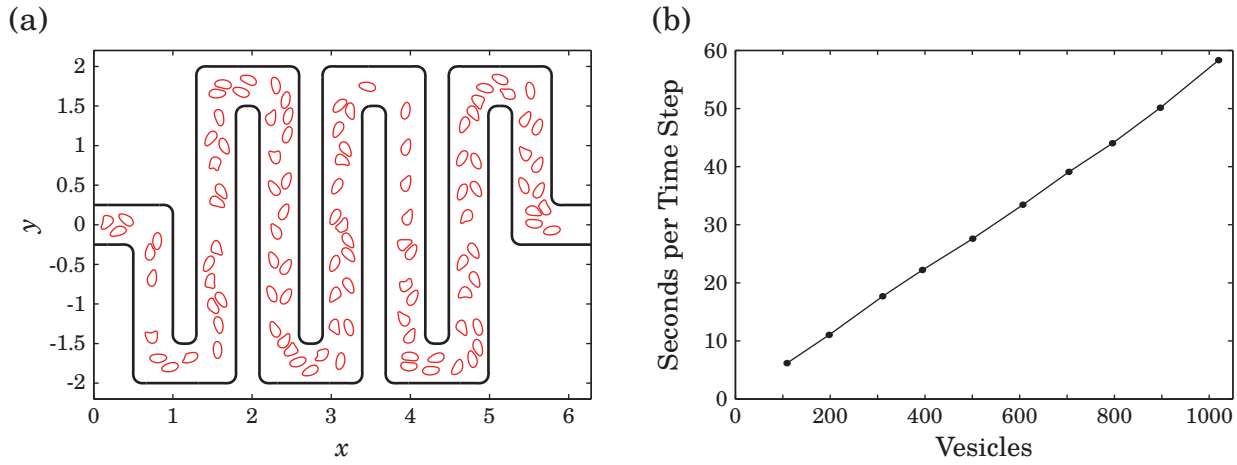

FIG. 5.5. *(a) A snapshot of 109 vesicles in the serpentine channel. A pressure difference is pushing the vesicles in the positive x-direction. (b) Average time per time step (in seconds) for the first 10 time steps as the number of vesicles is varied. Each dot represents a data point. Timings were performed using a laptop with a 2.4 GHz dual-core Intel Core i5 processor and 8 GB of RAM.*

time was used to construct a preconditioner which was a sparse block diagonal matrix consisting of vesicle-vesicle self interactions. Notice that while almost one half of the points are associated with the fixed geometry, the solve associated with these points takes less than 2% of the total time for the time step, thanks to the fast direct solver.

5.3. Fast direct solver. This section reports on the performance of the fast direct solver for the wall computations. The experiments in this section were performed on a laptop computer with two quad core Intel Core i7-4700MQ processors and 16 GB of RAM. The user prescribed tolerance was set to $\epsilon = 10^{-10}$.

Table 5.2 reports the time for the precomputation T_{pre}, the time for a solve T_{solve}, and the absolute error E for a point in the interior of the channel. For all geometries, the fast direct solver scales linearly. As the serpentine channel is under-discretized until $N = 4{,}000$, the precomputation step in the solver does not observe asymptotic complexity until $N = 32{,}000$. However, the solve step of the solver observes asymptotic complexity earlier.

Figure 5.6 reports the time in seconds for the precomputation (a) and the solve (b) steps verses the

TABLE 5.1

The CPU time distribution for the first 10 time steps of a simulation with 1,020 vesicles (64 points each) in the serpentine channel with $N = 29{,}580$ points per wall. Each time step took an average of 58 seconds on a laptop with a 2.4 GHz dual-core Intel Core i5 processor and 8 GB of RAM.

Operation	Percentage
vesicle to vesicle interactions	63.49
vesicle to wall interactions	14.77
preconditioner	7.30
wall to vesicle interactions	3.80
bending and tension forces	3.17
inextensibility operator	1.41
solve for τ, c using A^{-1}	1.40
vesicle to side interactions	1.31
proxy point to vesicle interactions	1.23
vesicle area corrections	1.16
other	0.96

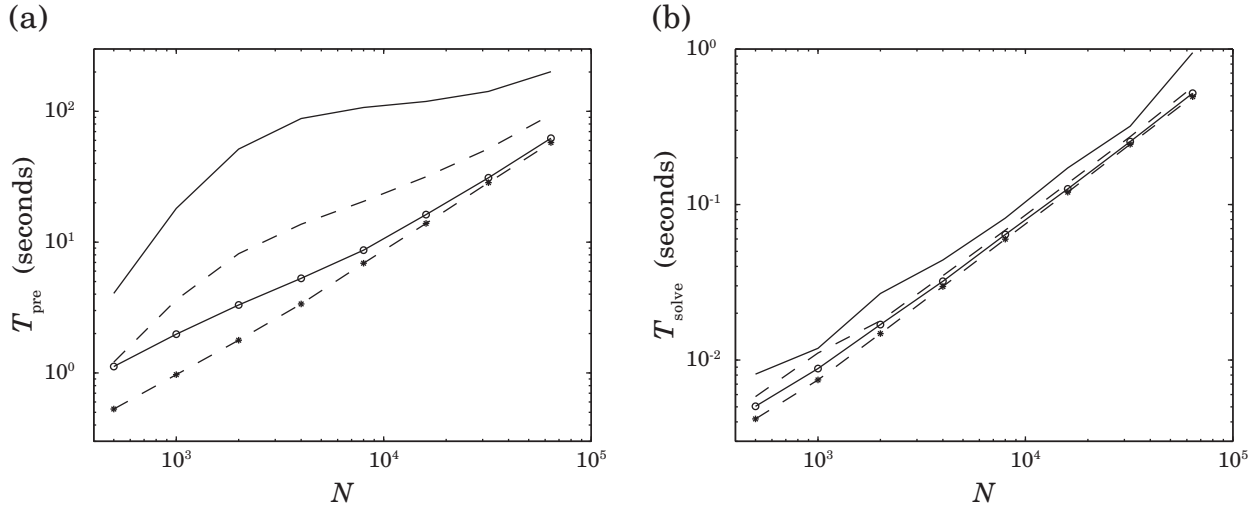

FIG. 5.6. (∗) *straight channel,* (∘) *converging-diverging channel,* (−) *serpentine channel,* (−−) *channel with reservoirs. Time in seconds for the precomputation (a) and solve (b) steps verses the number of discretization points N on the walls when the fast direct solver is applied to the geometries in Fig. 5.1.*

number of discretization points N on the walls for all of the geometries. The time for the precomputation of the straight channel and the channel with reservoirs is about the same since these geometries have approximately the same rank interactions. The rank interactions for the converging-diverging channel are smaller thus the constant for the precomputation is smaller. Thanks to the small constant associated with applying an HBS matrix, the cost for the solves is approximately the same independent of geometry.

In Table 5.3, we perform a comparison between the LU decomposition, GMRES, and the fast direct solver when computing $\boldsymbol{\tau}$ and $\mathbf{c}$ for the serpentine channel. In this example, we set $\epsilon = 10^{-12}$ for the fast direct solver. The timings (in seconds) do not include the initial LU factorization or the fast direct solver's compression of A^{-1}. Only the solve times are reported. We observe similar timings between the LU factorization and the fast direct solver for a relatively small number of points ($N \approx 400$). For larger values, the fast solver provides superior performance. In the comparison with GMRES, we used the FMM to apply A and reported the timings for the system without preconditioning. We found that GMRES required approximately 259 iterations to obtain a relative error of 10^{-12}. The number of iterations was independent of N and mainly depended on the complexity of the geometry. Notice that even a single GMRES iteration requires more time than the fast solver.

6. Conclusions. We presented a new algorithm for simulating particulate flows through arbitrary periodic geometries that is comprised mainly of three *independent* modules: a periodization scheme for Stokes flow through complex geometries, a fast free-space solver for simulating vesicle flows, and a fast direct solver for the boundary integral equations on the channel. We would like to emphasize that, owing to the linearity of Stokes equations, the periodization scheme can be combined seamlessly with any other existing free-space particulate flow solver implementations (e.g., those for bubbles, drops, rigid-particles, or swimmers). Similarly, other direct solver implementations can be used with equal ease. We showed via numerical experiments that the computational cost of the overall scheme scales linearly with respect to both space and time discretization sizes.

We are currently working on extending our work on several fronts. To enable higher volume-fraction flow simulations, we are developing a new chunk-based close evaluation scheme to accurately compute the channel-to-vesicle hydrodynamic interactions. In addition, the channel representation will support spatial adaptivity and the channel-vesicle interactions will be treated semi-implicitly. We plan to incorporate islands in the fluid domain by using the standard double-layer formulation as well as extending the scheme to three dimensions in the near future.

7. Acknowledgements. GM and SV acknowledge support from NSF under grants DMS-1224656, DMS-1418964 and DMS-1454010 and a Simons Collaboration Grant for Mathematicians No. 317933. AB

TABLE 5.2

Time in seconds for precomputation T_{pre} and solve T_{solve} as the number of discretization points N increases on the walls for the four geometries. The absolute error E at a point in the channel is also reported.

	Straight Channel			Converging-Diverging Channel		
N	T_{pre}	T_{solve}	E	T_{pre}	T_{solve}	E
500	0.53	4.18×10^{-3}	2.79×10^{-11}	1.12	5.04×10^{-3}	1.26×10^{-6}
1000	0.97	7.45×10^{-3}	1.10×10^{-9}	1.98	8.81×10^{-3}	1.12×10^{-9}
2000	1.78	1.48×10^{-2}	1.05×10^{-9}	3.31	1.69×10^{-2}	1.33×10^{-9}
4000	3.37	2.97×10^{-2}	6.05×10^{-9}	5.29	3.21×10^{-2}	7.28×10^{-9}
8000	6.9	5.99×10^{-2}	1.68×10^{-8}	8.68	6.41×10^{-2}	2.86×10^{-9}
16000	13.9	1.20×10^{-1}	1.13×10^{-8}	16.25	1.26×10^{-1}	1.00×10^{-8}
32000	28.5	2.44×10^{-1}	7.83×10^{-8}	31.06	2.54×10^{-1}	9.66×10^{-9}
64000	57.7	4.94×10^{-1}	1.92×10^{-7}	62.26	5.20×10^{-1}	8.28×10^{-9}

	Serpentine Channel			Channel with Reservoirs		
N	T_{pre}	T_{solve}	E	T_{pre}	T_{solve}	E
500	4.06	8.12×10^{-3}	8.62×10^{-2}	1.21	5.82×10^{-3}	1.53
1000	18.1	1.19×10^{-2}	2.72×10^{-2}	3.62	1.11×10^{-2}	3.20×10^{-3}
2000	51.5	2.68×10^{-2}	1.54×10^{-4}	8.16	1.79×10^{-2}	2.41×10^{-5}
4000	88.2	4.40×10^{-2}	3.18×10^{-8}	13.7	3.49×10^{-2}	1.22×10^{-8}
8000	107	8.18×10^{-2}	1.38×10^{-7}	20.5	6.84×10^{-2}	2.67×10^{-8}
16000	119	1.72×10^{-1}	2.29×10^{-9}	31.6	1.37×10^{-1}	3.35×10^{-8}
32000	142	3.19×10^{-1}	2.72×10^{-9}	51.8	2.73×10^{-1}	4.60×10^{-8}
64000	201	9.47×10^{-1}	1.10×10^{-8}	94.2	5.69×10^{-1}	6.80×10^{-8}

acknowledges support from NSF under grant DMS-1216656. This research was supported in part through computational resources and services provided by Advanced Research Computing at the University of Michigan, Ann Arbor.

TABLE 5.3

A comparison between the LU decomposition, GMRES, and the fast direct solver when computing τ and c for the serpentine channel. The timings (in seconds) do not include the initial LU factorization or the fast direct solver's compression of A^{-1}.

N	LU	GMRES	Fast Apply
500	1.87×10^{-2}	23.0	1.11×10^{-2}
1000	5.38×10^{-2}	42.3	2.29×10^{-2}
2000	1.89×10^{-1}	64.4	4.40×10^{-2}
4000	5.17×10^{-1}	91.1	9.06×10^{-2}
8000	–	158	1.64×10^{-1}
16000	–	267	3.70×10^{-1}

Appendix A. Arc length correction. In this section, we present a derivation of the arc length correction formula. The arc length correction is useful for long-time vesicle simulations where the accumulation of errors in the arc length may become significant, leading to elongated vesicles. By placing a small correction term on the right-hand side of the inextensiblity condition, we prevent this accumulation while preserving the membrane's original length with second-order asymptotic convergence.

We begin by understanding how errors accumulate. Let $\mathbf{x}^k(\alpha)$, where $\alpha \in [0, 2\pi)$, be a parameterization of the membrane γ_k. The cumulative error on the kth time step, denoted by $\mathcal{E}_k$, is given by

$$\mathcal{E}_k = \int_{\gamma_k} ds - \int_{\gamma_0} ds. \tag{A.1}$$

To find $\mathcal{E}_{k+1}$, we will use the identity

$$\begin{aligned}
\|\mathbf{x}_\alpha^{k+1}\|_2 &= \|\mathbf{x}_\alpha^k + \Delta t \mathbf{u}_\alpha\|_2 && \text{(A.2)}\\
&= \left(\left(x_\alpha^k + \Delta t u_\alpha\right)^2 + \left(y_\alpha^k + \Delta t v_\alpha\right)^2\right)^{\frac{1}{2}} && \text{(A.3)}\\
&= \left(\left((x_\alpha^k)^2 + (y_\alpha^k)^2\right) + 2\Delta t\left(u_\alpha x_\alpha^k + v_\alpha y_\alpha^k\right) + \Delta t^2\left((u_\alpha)^2 + (v_\alpha)^2\right)\right)^{\frac{1}{2}} && \text{(A.4)}\\
&= \left(1 + 2\Delta t(\mathbf{x}_s^k \cdot \mathbf{u}_s) + \Delta t^2 \|\mathbf{u}_s\|_2^2\right)^{\frac{1}{2}} \|\mathbf{x}_\alpha^k\|_2, && \text{(A.5)}
\end{aligned}$$

where $\mathbf{x}^k = (x^k, y^k)$ and $\mathbf{u} = (u, v)$. The error is then

$$\mathcal{E}_{k+1} = \int_{\gamma^k} \left(1 + 2\Delta t(\mathbf{x}_s^k \cdot \mathbf{u}_s) + \Delta t^2 \|\mathbf{u}_s\|_2^2\right)^{\frac{1}{2}} ds - \int_{\gamma_0} ds. \tag{A.6}$$

Observe that setting $\mathbf{x}_s^k \cdot \mathbf{u}_s = 0$ will cause the arc length to increase monotonically since $\|\mathbf{u}_s\|_2^2 \geq 0$. We now perform a Taylor expansion around $\Delta t = 0$ to get

$$\mathcal{E}_{k+1} = \mathcal{E}_k + \Delta t \int_{\gamma^k} \mathbf{x}_s^k \cdot \mathbf{u}_s \, ds + \frac{1}{2}\Delta t^2 \int_{\gamma^k} \|\mathbf{u}_s\|_2^2 \, ds - \frac{1}{2}\Delta t^2 \int_{\gamma^k} (\mathbf{x}_s^k \cdot \mathbf{u}_s)^2 \, ds + \mathcal{O}(\Delta t^3). \tag{A.7}$$

If we set $\mathbf{x}_s^k \cdot \mathbf{u}_s = 0$, we see that with each time step, the arc length incurs an error on the order of Δt^2. If n is the total number of time steps, we would then expect the cumulative error to scale like $\mathcal{E}_n = \mathcal{O}(n\Delta t^2)$. This poses a problem for long-time simulations, where n is very large. To address this, we need to prevent the $\mathcal{E}_k$ term in (A.7) from causing any significant accumulation. A simple remedy is to let

$$\mathbf{x}_s^k \cdot \mathbf{u}_s = -\frac{\mathcal{E}_k}{\Delta t \int_{\gamma^k} ds} = \frac{\int_{\gamma^0} ds - \int_{\gamma^k} ds}{\Delta t \int_{\gamma^k} ds}. \tag{A.8}$$

Each time step, we are still incurring an error on the order of Δt^2. However, we propose that the cumulative error now scales $\mathcal{E}_n = \mathcal{O}(\Delta t^2)$ and that there exists an upper bound for the relative error that scales Δt^2, independent of n. We only require that $\|\mathbf{u}_s\|_2$ be bounded.

To begin, let $L_k = \int_{\gamma^k} ds$. Using (A.1), (A.6), and (A.8), we find that

$$L_{k+1} = \int_{\gamma^k} \left(1 + 2\frac{L_0 - L_k}{L_k} + \Delta t^2 \|\mathbf{u}_s\|_2^2\right)^{\frac{1}{2}} ds. \tag{A.9}$$

LEMMA A.1. *Suppose there exists a constant C such that $\|\mathbf{u}_s\|_2 \leq C$ for all time. Then,*

$$1 - \frac{\Delta t^2 C^2}{2 - \Delta t^2 C^2} \leq \frac{L_k}{L_0} \leq 1 + \frac{\Delta t^2 C^2}{2 - \Delta t^2 C^2} \tag{A.10}$$

for all $k \geq 0$ when Δt is sufficiently small.

Proof. Clearly, (A.10) holds for the base case $L_k = L_0$. Assume the induction hypothesis. We will use the fact that

$$\frac{L_k}{L_0}\left(1+2\frac{L_0-L_k}{L_k}\right)^{\frac{1}{2}} \leq \frac{L_{k+1}}{L_0} \leq \frac{L_k}{L_0}\left(1+2\frac{L_0-L_k}{L_k}+\Delta t^2 C^2\right)^{\frac{1}{2}} \tag{A.11}$$

to place the desired bounds on L_{k+1}/L_0.

We begin by showing

$$1-\frac{\Delta t^2 C^2}{2-\Delta t^2 C^2} \leq \frac{L_k}{L_0}\left(1+2\frac{L_0-L_k}{L_k}\right)^{\frac{1}{2}}. \tag{A.12}$$

For simplicity, let $z = L_k/L_0$ and $r = \Delta t^2 C^2$. We need to show that

$$1-\frac{r}{2-r} \leq z\left(\frac{2}{z}-1\right)^{\frac{1}{2}}, \quad \text{for} \quad z \in I_r = \left[1-\frac{r}{2-r}, 1+\frac{r}{2-r}\right]. \tag{A.13}$$

The right-hand side of the inequality in (A.13) is the top half of a circle centered about the point $(1,0)$ with radius 1. The minimum occurs on both endpoints of I_r and is given by

$$\min_{z\in I_r} z\left(\frac{2}{z}-1\right)^{\frac{1}{2}} = \frac{2}{2-r}\left(1-r\right)^{\frac{1}{2}}. \tag{A.14}$$

We find that (A.13) holds as long as $r \leq 1$. Therefore, (A.12) holds as long as $\Delta t \leq 1/C$.

All that remains is to show that

$$\frac{L_k}{L_0}\left(1+2\frac{L_0-L_k}{L_k}+\Delta t^2 C^2\right)^{\frac{1}{2}} \leq 1+\frac{\Delta t^2 C^2}{2-\Delta t^2 C^2}, \tag{A.15}$$

which is equivalent to

$$z\left(\frac{2}{z}-1+r\right)^{\frac{1}{2}} \leq 1+\frac{r}{2-r}, \quad \text{for} \quad z \in I_r. \tag{A.16}$$

The left-hand side of the inequality in (A.16) is the top half of an ellipse centered about $(\frac{1}{1-r}, 0)$. The maximum occurs on the right endpoint of I_r and is given by

$$\max_{z\in I_r} z\left(\frac{2}{z}-1+r\right)^{\frac{1}{2}} = 1+\frac{r}{2-r}. \tag{A.17}$$

Therefore, (A.16) holds, which completes the proof of Lemma A.1. □

THEOREM A.2. *Suppose there exists a constant C such that $\|\mathbf{u}_s\|_2 \leq C$ for all time. Then, the condition*

$$\mathbf{x}_s^k \cdot \mathbf{u}_s = \frac{L_0-L_k}{\Delta t L_k}, \tag{A.18}$$

preserves arc length with relative error

$$\left|\frac{L_{k+1}-L_0}{L_0}\right| \leq \frac{\Delta t^2 C^2}{2-\Delta t^2 C^2} = \frac{\Delta t^2 C^2}{2} + \mathcal{O}(\Delta t^4) \tag{A.19}$$

for all $k \geq 0$ when $\Delta t \leq 1/C$.

TABLE A.1

Convergence analysis for the arc length and area corrections. To perform the analysis, we evolved a vesicle using time step $\Delta t = 0.01/M$ until $t = 0.1$. We report the relative errors in the arc length L, area A, and position $\mathbf{x}$ with and without the corrections. The subscript c indicates that both the area and arc length corrections were used, and the subscript nc indicates that no corrections were used. The accepted values are labeled with subscript acc. We now observe second-order asymptotic convergence for the arc length, which is in agreement with Theorem A.2.

M	$\left\lvert\frac{L_c - L_{acc}}{L_{acc}}\right\rvert$	$\left\lvert\frac{A_c - A_{acc}}{A_{acc}}\right\rvert$	$\left\lVert\frac{\mathbf{x}_c - \mathbf{x}_{acc}}{L_{acc}}\right\rVert_\infty$	$\left\lvert\frac{L_{nc} - L_{acc}}{L_{acc}}\right\rvert$	$\left\lvert\frac{A_{nc} - A_{acc}}{A_{acc}}\right\rvert$	$\left\lVert\frac{\mathbf{x}_{nc} - \mathbf{x}_{acc}}{L_{acc}}\right\rVert_\infty$
1	8.65×10^{-5}	1.21×10^{-10}	3.02×10^{-4}	1.15×10^{-3}	3.12×10^{-4}	4.78×10^{-4}
2	2.09×10^{-5}	7.16×10^{-12}	1.55×10^{-4}	5.77×10^{-4}	1.57×10^{-4}	2.38×10^{-4}
4	5.15×10^{-6}	4.35×10^{-13}	7.80×10^{-5}	2.90×10^{-4}	7.91×10^{-5}	1.17×10^{-4}
8	1.28×10^{-6}	2.72×10^{-14}	3.90×10^{-5}	1.45×10^{-4}	3.97×10^{-5}	5.68×10^{-5}
16	3.18×10^{-7}	3.09×10^{-15}	1.93×10^{-5}	7.26×10^{-5}	1.99×10^{-5}	2.65×10^{-5}

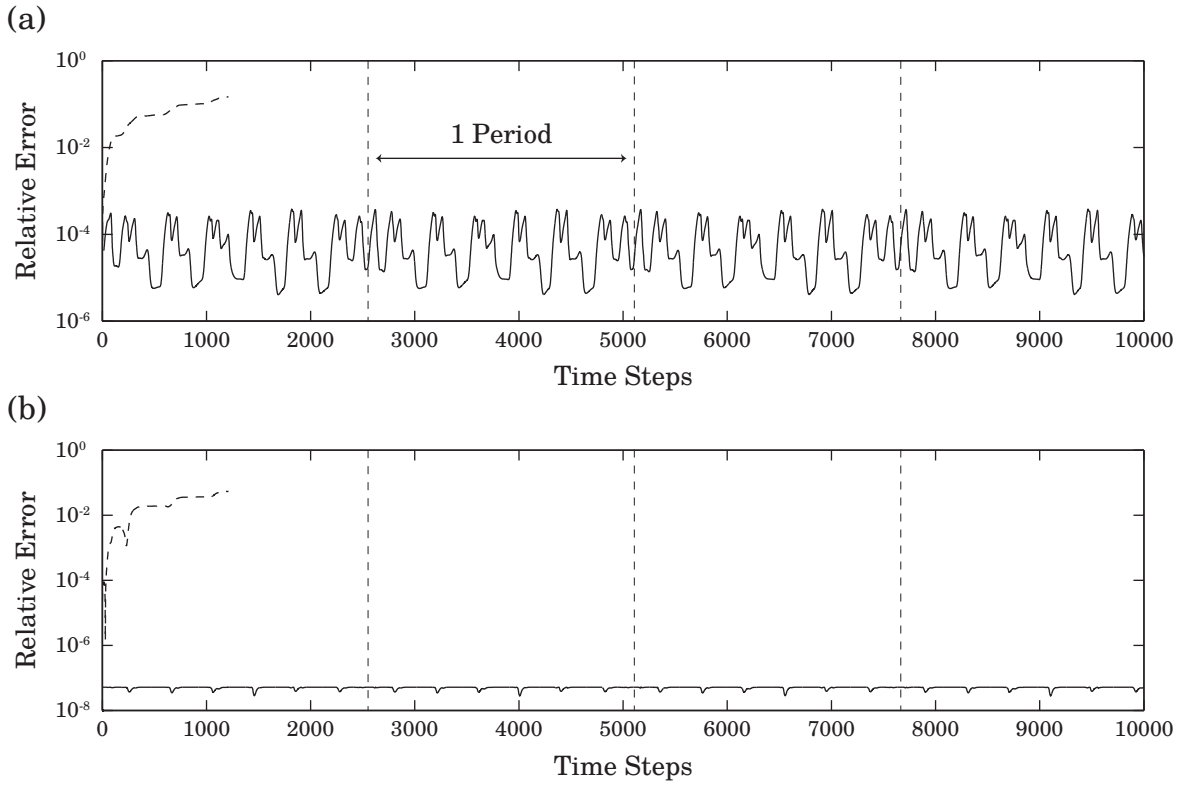

FIG. A.1. *Errors without corrections (−−) and with both arc length and area corrections (−). The relative error in arc length (a) and area (b) for a single vesicle flowing through the serpentine channel. After around $\sim 1,200$ time steps, the simulation breaks down due to numerical instabilities when no corrections are employed. In the simulation using both corrections, the vesicle successfully passed through the channel several times without incident. In both simulations, the time step was set to $\Delta t = 0.01$.*

Appendix B. Area correction. Although the fluid incompressibility condition is satisfied exactly by the single-layer kernel, errors in the area of vesicles (owing to discretization) accumulate over time and have a compounding effect in the case of long-time simulations. In this section, we discuss a simple and efficient method to correct the area errors whenever the vesicle shapes are updated i.e., at every time step. Let $\mathbf{x}^k(\alpha) = (x^k(\alpha), y^k(\alpha))$ with $\alpha \in [0, 2\pi)$ represent the position of a vesicle's membrane on the

kth time step. The initial area is given by

$$A_0 = \int_0^{2\pi} x^0 y^0_\alpha \, d\alpha. \tag{B.1}$$

To correct $\mathbf{x}^k$, we simply add a normal vector $\mathbf{n} = (y^k_\alpha, -x^k_\alpha)$ scaled by a small unknown constant c, which is computed by requiring that the area enclosed by $\mathbf{x}^k + c\mathbf{n}$ equals A_0. That is,

$$\int_0^{2\pi} \left(x^k + cy^k_\alpha\right)\left(y^k - cx^k_\alpha\right)_\alpha \, d\alpha = A_0. \tag{B.2}$$

Expanding the integrand gives

$$\int_0^{2\pi} x^k y^k_\alpha \, d\alpha - c\int_0^{2\pi} x x^k_{\alpha\alpha} \, d\alpha + c\int_0^{2\pi} \left(y^k_\alpha\right)^2 \, d\alpha - c^2 \int_0^{2\pi} y^k_\alpha x^k_{\alpha\alpha} \, d\alpha = A_0. \tag{B.3}$$

We take c to be the closest root to zero of this quadratic equation.

Appendix C. Fast reparameterization. The classical approach for reparameterizing evolving geometries in 2D is to introduce an auxiliary tangential velocity in the kinematic condition that helps maintain parameterization quality under some metric (e.g., equispaced in arclength) [28]. This approach is very effective and widely used. In our setting, however, in addition to evolving the shape, we need to keep track of the membrane tension from previous time-step. While this can be accomplished by advecting the scalar field (tension) with the auxiliary velocity field, it effects the overall accuracy and stability of our time-stepping scheme. Moreover, our requirements are different: we do not need to maintain exact equi-arclength parameterization but rather a scheme that overcomes the error-amplification due to large shape deformations. For this reason, we take a different approach that avoids an auxiliary advection equation solve at each time step. The scheme only uses Fourier interpolation and redistributes the discretization points on the vesicle membrane so that they are *nearly* equispaced.

Let $\mathbf{x}(\alpha)$, where $\alpha \in [0, 2\pi)$, be a parameterization for a vesicle's membrane γ. We begin by expressing $ds/d\alpha$ as a Fourier series. That is,

$$\frac{ds}{d\alpha} = \|\mathbf{x}_\alpha\|_2 = \sum_{k=-M/2+1}^{M/2} C_k e^{ik\alpha}, \tag{C.1}$$

where M is the number of discretization points. Integrating both sides gives

$$s(\alpha) = \int_0^\alpha \|\mathbf{x}_\beta\|_2 \, d\beta = K + C_0\alpha + \sum_{k\neq 0} \frac{C_k}{ik} e^{ik\alpha}, \tag{C.2}$$

where

$$K = -\sum_{k\neq 0} \frac{C_k}{ik}. \tag{C.3}$$

Our goal is to find a set of discrete points $\{\alpha^*_k\}_{k=0}^{M-1}$ in the parametric domain that corresponds to $\{s_j = s(2\pi j/M), j = 0, \ldots, M-1\}$. We simply use a piecewise linear interpolant of $s(\alpha)$ to determine these points; see Algorithm 2 for more details. The result is that we obtain points on the curve that are nearly equispaced (up to the accuracy of linear interpolation). The coordinate positions $\mathbf{x}$ are then evaluated at α^* via Fourier interpolation. Note that the linear interpolation does not interfere with the overall spectral accuracy of the method (in space) because it is used only to find a new set of points α^*.

REFERENCES

[1] L. Af Klinteberg and A.-K. Tornberg, *Fast Ewald summation for Stokesian particle suspensions*, International Journal for Numerical Methods in Fluids, 76 (2014), pp. 669–698.

Algorithm 2 Compute α^*

$\alpha_0^* = 0$, $j = 0$
for $k = 1 : M-1$ **do**
  while $\frac{s(2\pi)}{M} k \notin [s_j, s_{j+1})$ **do**
    $j = j + 1$
  end while
  $\alpha_k^* = \frac{2\pi}{M}\left(j + \frac{\frac{s(2\pi)}{M}k - s_j}{s_{j+1} - s_j}\right)$
end for

[2] S. Ambikasaran and E. Darve, *An $o(n \log n)$ fast direct solver for partial hierarchically semi-separable matrices*, J. Sci. Comput., 57 (2013), pp. 477–501.

[3] U. M. Ascher and L. R. Petzold, *Computer Methods for Ordinary Differential Equations and Differential-Algebraic Equations*, Society for Industrial and Applied Mathematics, Philadelphia, PA, USA, 1998.

[4] A. H. Barnett and T. Betcke, *Stability and convergence of the Method of Fundamental Solutions for Helmholtz problems on analytic domains*, J. Comput. Phys., 227 (2008), pp. 7003–7026.

[5] A. H. Barnett and L. Greengard, *A new integral representation for quasi-periodic scattering problems in two dimensions*, BIT Numer. Math., 51 (2011), pp. 67–90.

[6] A. H. Barnett, B. Wu, and S. Veerapaneni, *Spectrally-accurate quadratures for evaluation of layer potentials close to the boundary for the 2D Stokes and Laplace equations*, SIAM J. Sci. Comput., 37 (2015), pp. B519–B542.

[7] A. Bogomolny, *Fundamental solutions method for elliptic boundary value problems*, SIAM J. Numer. Anal., 22 (1985), pp. 644–669.

[8] S. Börm, *Efficient numerical methods for non-local operators*, vol. 14 of EMS Tracts in Mathematics, European Mathematical Society (EMS), Zürich, 2010. $\mathcal{H}^2$-matrix compression, algorithms and analysis.

[9] S. Börm and W. Hackbusch, *Approximation of boundary element operators by adaptive $\mathcal{H}^2$-matrices*, in Foundations of computational mathematics: Minneapolis, 2002, vol. 312 of London Math. Soc. Lecture Note Ser., Cambridge Univ. Press, Cambridge, 2004, pp. 58–75.

[10] S. Chandrasekaran and M. Gu, *A divide-and-conquer algorithm for the eigendecomposition of symmetric block-diagonal plus semiseparable matrices*, Numer. Math., 96 (2004), pp. 723–731.

[11] H. Cheng, Z. Gimbutas, P. Martinsson, and V. Rokhlin, *On the compression of low rank matrices*, SIAM Journal of Scientific Computing, 26 (2005), pp. 1389–1404.

[12] M. H. Cho and A. H. Barnett, *Robust fast direct integral equation solver for quasi-periodic scattering problems with a large number of layers*, Opt. Express, 23 (2015), pp. 1775–1799.

[13] D. Colton and R. Kress, *Inverse acoustic and electromagnetic scattering theory*, vol. 93 of Applied Mathematical Sciences, Springer-Verlag, Berlin, second ed., 1998.

[14] T. Darden, D. York, and L. Pedersen, *Particle mesh Ewald: An $N \log(N)$ method for Ewald sums in large systems*, J. Chem. Phys., 98 (1993), p. 10089.

[15] P. Decuzzi, B. Godin, T. Tanaka, S.-Y. Lee, C. Chiappini, X. Liu, and M. Ferrari, *Size and shape effects in the biodistribution of intravascularly injected particles.*, Journal of controlled release : official journal of the Controlled Release Society, 141 (2010), pp. 320–7.

[16] P. P. Ewald, *Die berechnung optischer und elektrostatischer gitterpotentiale*, Annalen der Physik, 369 (1921), pp. 253–287.

[17] X.-J. Fan, N. Phan-Thien, and R. Zheng, *Completed double layer boundary element method for periodic suspensions*, Zeitschrift für angewandte Mathematik und Physik ZAMP, 49 (1998), pp. 167–193.

[18] J. Freund, *Leukocyte margination in a model microvessel*, Physics of Fluids, 19 (2007), p. 023301.

[19] A. Gholami, D. Malhotra, H. Sundar, and G. Biros, *Fft, fmm, or multigrid? a comparative study of state-of-the-art poisson solvers*, arXiv preprint arXiv:1408.6497, (2014).

[20] A. Gillman and A. Barnett, *A fast direct solver for quasi-periodic scattering problems*, J. Comput. Phys., 248 (2013), pp. 309–322.

[21] A. Gillman, P. Young, and P. Martinsson, *A direct solver with $O(N)$ complexity for integral equations on one-dimensional domains*, Frontiers of Mathematics in China, 7 (2012), pp. 217–247.

[22] Z. Gimbutas and L. Greengard, *STFMMLIB3 - Fast Multipole Method (FMM) library for the evaluation of potential fields governed by the Stokes equations in R^3*. `http://www.cims.nyu.edu/cmcl/fmm3dlib/fmm3dlib.html`, 2012.

[23] L. Greengard and M. C. Kropinski, *Integral equation methods for stokes flow in doubly-periodic domains*, Journal of engineering mathematics, 48 (2004), pp. 157–170.

[24] M. Gu and S. C. Eisenstat, *Efficient algorithms for computing a strong rank-revealing QR factorization*, SIAM J. Sci. Comput., 17 (1996), pp. 848–869.

[25] H. Hasimoto, *On the periodic fundamental solutions of the Stokes equations and their application to viscous flow past a cubic array of spheres*, Journal of Fluid Mechanics, 5 (1959), pp. 317–328.

[26] W. Helfrich, *Elastic properties of lipid bilayers: theory and possible experiments*, Zeitschrift für Naturforschung C, 28 (1973), pp. 693–703.

[27] J. Helsing and R. Ojala, *On the evaluation of layer potentials close to their sources*, J. Comput. Phys., 227 (2008), pp. 2899–2921.

[28] T. Y. Hou, J. S. Lowengrub, and M. J. Shelley, *Removing the stiffness from interfacial flows with surface tension*, Journal of Computational Physics, 114 (1994), pp. 312–338.
[29] G. Hsiao and W. L. Wendland, *Boundary Integral Equations*, Applied Mathematical Sciences, Vol. 164, Springer, 2008.
[30] R. B. Huang, S. Mocherla, M. J. Heslinga, P. Charoenphol, and O. Eniola-Adefeso, *Dynamic and cellular interactions of nanoparticles in vascular-targeted drug delivery (review).*, Molecular membrane biology, 27 (2010), pp. 190–205.
[31] C. Jin, S. M. McFaul, S. P. Duffy, X. Deng, P. Tavassoli, P. C. Black, and H. Ma, *Technologies for label-free separation of circulating tumor cells: from historical foundations to recent developments*, Lab on a Chip, 14 (2014), pp. 32–44.
[32] Y. G. K. and A. Acrivos, *Stokes flow past a particle of arbitrary shape: a numerical method of solution*, Journal of Fluid Mechanics, 69 (1975), pp. 377–403.
[33] ———, *On the shape of a gas bubble in a viscous extensional flow*, Journal of Fluid Mechanics, 76 (1976), pp. 433–442.
[34] R. Kress, *Boundary integral equations in time-harmonic acoustic scattering*, Mathl. Comput. Modelling, 15 (1991), pp. 229–243.
[35] ———, *Numerical Analysis*, Graduate Texts in Mathematics #181, Springer-Verlag, 1998.
[36] R. Kress, *Linear Integral Equations*, vol. 82 of Appl. Math. Sci., Springer, second ed., 1999.
[37] O. A. Ladyzhenskaya, *The Mathematical Theory of Viscous Incompressible Flow, revised 2nd edition*, Mathematics and Its Applications 2, Gordon and Breach, 1969.
[38] R. Larson and J. Higdon, *Microscopic flow near the surface of two-dimensional porous media. part 1. axial flow*, Journal of Fluid Mechanics, 166 (1986), pp. 449–472.
[39] D. Lindbo and A.-K. Tornberg, *Spectrally accurate fast summation for periodic Stokes potentials*, Journal of Computational Physics, 229 (2010), pp. 8994–9010.
[40] Y. Liu, *Fast Multipole Boundary Element Method: Theory and Applications in Engineering*, Cambridge University Press, 2009.
[41] Y. Liu and A. Barnett, *Efficient numerical solution of acoustic scattering from doubly-periodic arrays of axisymmetric objects*, 2015. submitted, *J. Comput. Phys.*; `arxiv:1506.05083`.
[42] M. Loewenberg and E. Hinch, *Numerical simulations of concentrated emulsions*, Journal of Fluid Mechanics, 321 (1996), pp. 395–419.
[43] R. Ojala and A.-K. Tornberg, *An accurate integral equation method for simulating multi-phase Stokes flow*, arXiv preprint arXiv:1404.3552, (2014).
[44] W. Olbricht, *Pore-scale prototypes of multiphase flow in porous media*, Annual review of fluid mechanics, 28 (1996), pp. 187–213.
[45] C. Pozrikidis, *Boundary integral and sinularity methods for linearized viscous flow*, Cambridge University Press, 1992.
[46] C. Pozrikidis, *Computation of periodic Green's functions of Stokes flow*, Journal of engineering Mathematics, 30 (1996), pp. 79–96.
[47] C. Pozrikidis, *Interfacial dynamics for Stokes flow*, Journal of Computational Physics, 169 (2001), pp. 250–301.
[48] A. Rahimian, S. K. Veerapaneni, and G. Biros, *Dynamic simulation of locally inextensible vesicles suspended in an arbitrary two-dimensional domain, a boundary integral method*, Journal of Computational Physics, 229 (2010), pp. 6466–6484.
[49] A. Russom, A. K. Gupta, S. Nagrath, D. Di Carlo, J. F. Edd, and M. Toner, *Differential inertial focusing of particles in curved low-aspect-ratio microchannels.*, New journal of physics, 11 (2009), p. 75025.
[50] Z. Sheng, P. Dewilde, and S. Chandrasekaran, *Algorithms to solve hierarchically semi-separable systems*, in System theory, the Schur algorithm and multidimensional analysis, vol. 176 of Oper. Theory Adv. Appl., Birkhäuser, Basel, 2007, pp. 255–294.
[51] M. Thiébaud and C. Misbah, *Rheology of a vesicle suspension with finite concentration: A numerical study*, Physical Review E, 88 (2013), p. 062707.
[52] A.-K. Tornberg and M. Shelley, *Simulating the dynamics and interactions of flexible fibers in Stokes flows*, Journal of Computational Physics, 196 (2004), pp. 8–40.
[53] S. K. Veerapaneni, D. Gueyffier, D. Zorin, and G. Biros, *A boundary integral method for simulating the dynamics of inextensible vesicles suspended in a viscous fluid in 2d*, J. Comput. Phys., 228 (2009), pp. 2334–2353.
[54] P. M. Vlahovska, D. Barthes-Biesel, and C. Misbah, *Flow dynamics of red blood cells and their biomimetic counterparts*, Comptes Rendus Physique, 14 (2013), pp. 451–458.
[55] M. Wang and J. F. Brady, *Spectral Ewald acceleration of Stokesian dynamics for polydisperse suspensions*, arXiv preprint arXiv:1506.08481, (2015).
[56] J. Xia, S. Chandrasekaran, M. Gu, and X. Li, *Fast algorithms for hierarchically semiseparable matrices*, Numer. Linear Algebra Appl., 17 (2010), pp. 953–976.
[57] J. Xia, S. Chandrasekaran, M. Gu, and X. S. Li, *Superfast multifrontal method for large structured linear systems of equations*, SIAM J. Matrix Anal. Appl., 31 (2009), pp. 1382–1411.
[58] H. Zhao, A. H. G. Isfahani, L. N. Olson, and J. B. Freund, *A spectral boundary integral method for flowing blood cells*, J. Comput. Phys., 229 (2010), pp. 3726–3744.
[59] H. Zhao and E. S. Shaqfeh, *The dynamics of a non-dilute vesicle suspension in a simple shear flow*, Journal of Fluid Mechanics, 725 (2013), pp. 709–731.
[60] A. Zick and G. Homsy, *Stokes flow through periodic arrays of spheres*, Journal of fluid mechanics, 115 (1982), pp. 13–26.
[61] A. Zinchenko and R. Davis, *An Efficient Algorithm for Hydrodynamical Interaction of Many Deformable Drops*,

Journal of Computational Physics, 157 (2000), pp. 539–587.